

A lunar radio experiment with the Parkes radio telescope for the LUNASKA project

J.D. Bray^{a,b,c,*}, R.D. Ekers^b, P. Roberts^b, J.E. Reynolds^b, C.W. James^d, C.J. Phillips^b, R.J. Protheroe^a, R.A. McFadden^e, M.G. Aartsen^a

^a*School of Chemistry & Physics, Univ. of Adelaide, SA 5005, Australia*

^b*CSIRO Astronomy & Space Science, Epping, NSW 1710, Australia*

^c*School of Physics & Astronomy, Univ. of Southampton, SO17 1BJ, United Kingdom*

^d*ECAP, Univ. of Erlangen-Nuremberg, 91058 Erlangen, Germany*

^e*ASTRON, 7990 AA Dwingeloo, The Netherlands*

Abstract

We describe an experiment using the Parkes radio telescope in the 1.2–1.5 GHz frequency range as part of the LUNASKA project, to search for nanosecond-scale pulses from particle cascades in the Moon, which may be triggered by ultra-high-energy astroparticles. Through the combination of a highly sensitive multi-beam radio receiver, a purpose-built backend and sophisticated signal-processing techniques, we achieve sensitivity to radio pulses with a threshold electric field strength of $0.0053 \mu\text{V}/\text{m}/\text{MHz}$, lower than previous experiments by a factor of three. We observe no pulses in excess of this threshold in observations with an effective duration of 127 hours. The techniques we employ, including compensating for the phase, dispersion and spectrum of the expected pulse, are relevant for future lunar radio experiments.

Keywords: ultra-high-energy neutrinos, ultra-high-energy cosmic rays, lunar radio experiment, radio instrumentation

1. Introduction

Studies of ultra-high-energy (UHE; $> 10^{18}$ eV) cosmic rays, and searches for their expected counterpart neutrinos, are difficult because of their extremely low flux. Current observatories such as the Pierre Auger Observatory [1] and the Telescope Array [2] make use of arrays of detectors covering thousands of km^2 to detect the particle cascades produced when UHE particles interact in the Earth's atmosphere. An alternative approach [3] is to use the Moon as the detector, searching for the nanosecond-scale radio pulse produced when a UHE particle interacts in a dense medium such as the lunar regolith [4] with terrestrial radio telescopes. This was originally proposed as a means to detect UHE neutrinos, but is also sensitive to cosmic rays [5–7], and the large size of the Moon gives it a larger potential aperture than other detection techniques. The detection of an individual nanosecond-scale pulse requires techniques that are quite unlike those of conventional radio astronomy, and are being refined through ongoing experimentation, with the ultimate goal of establishing lunar radio observations as a practical technique for detecting UHE particles.

The first lunar particle detection experiment was performed in 1995 with the Parkes radio telescope [8]; subsequent experiments have been carried out with the Goldstone Deep Space Communications Complex (GDSCC) [9, GLUE], the Kalyazin radio telescope [10], the Australia

Telescope Compact Array (ATCA) [11, LUNASKA], the Lovell telescope [12, La Luna], the Westerbork Synthesis Radio Telescope (WSRT) [13, NuMoon] and the Expanded Very Large Array (EVLA) [14, RESUN]. A key feature of these experiments is that they make use of existing radio telescopes, whereas other experiments searching for UHE cosmic rays or neutrinos require the development of expensive dedicated instruments. The sensitivity of radio telescopes is undergoing rapid improvement, through the upgrading of existing instruments and the construction of new ones, which can be exploited to achieve greater sensitivity to UHE particles. These efforts are expected to culminate in the Square Kilometre Array (SKA) [15], an international radio telescope scheduled to begin construction in 2016, which will be sensitive even to the minimal neutrino flux expected from the Greisen-Zatsepin-Kuzmin [16, 17] interactions of propagating cosmic rays [5]. The LUNASKA (Lunar Ultra-high energy Neutrino Astrophysics with the SKA) project aims to develop the lunar radio approach to particle detection for future use with the SKA.

Our previous LUNASKA experiment searched for coincident pulses independently detected by three antennas of the ATCA, and was therefore limited by the sensitivity of a single 22 m antenna. Improving on this with the same instrument would require combining the signals from multiple antennas into tied-array beams, similar to the NuMoon experiment with the WSRT, which would be prohibitively difficult given the ATCA's higher bandwidth and angular resolution. Instead, we conducted a further experiment with the 64 m Parkes radio telescope, with its greater col-

*Corresponding author.

Email address: j.bray@soton.ac.uk (J.D. Bray)

lecting area allowing an improvement in sensitivity. The disadvantage of this larger antenna is that its beam observes a smaller fraction of the Moon, which we counteracted by using a multibeam receiver to observe multiple regions of the Moon simultaneously as suggested by Ekers et al. [18]. The use of a single antenna also required a different strategy from the ATCA experiment for excluding pulses from radio frequency interference (RFI), which were identified primarily by their simultaneous appearance in multiple beams.

In this article, we describe our lunar radio experiment with the Parkes radio telescope in 2010 as part of the LUNASKA project, with improved sensitivity over our previous experiment with the ATCA. Sec. 2 contains an overview of the design of the experiment. In Secs. 3, 4 and 5 we go into more detail respectively about the calibration procedure, the effects of ionospheric dispersion, and the optimisation of the signal-to-noise ratio. In Sec. 6 we describe the analysis of the data from our observations, including the identification and exclusion of RFI, and in Sec. 7 we conclude by summarising our results and considerations for future experiments. The limits that our experiment establishes on the ultra-high-energy particle flux will be addressed in a separate work [19].

2. Experiment design

The Parkes radio telescope ($33^{\circ} 00' \text{ S } 148^{\circ} 16' \text{ E}$) is located in New South Wales, Australia, and consists of a single 64 m parabolic antenna with receivers mounted at the prime focus. This experiment used the 21 cm multibeam receiver [20], which is capable of observing up to thirteen points on the sky simultaneously. At any one time we used four of the receiver's thirteen beams, each with two orthogonal linear polarisations; see Sec. 2.1 for details of our pointing strategy. The receiver has a radio frequency (RF) band of 1.2–1.5 GHz, which is downconverted to an intermediate frequency (IF) band of 50–350 MHz.

Further processing of the signal was performed with Bedlam, a digital backend developed specifically for this experiment [21]. The voltage waveform was digitised with 8 bits of precision at a rate of 1,024 Msamples/s, oversampling the band by a factor ~ 1.7 . A copy of this digital signal was then passed through a digital filter to compensate for ionospheric dispersion (see Secs. 4 and 5.2); both the raw and dedispersed data were stored in temporary buffers, with an adjustable size of up to 8 μs . Fig. 1 shows details of the signal path.

Due to the high data rate, the buffered data were copied to permanent storage only if triggered by a set of conditions corresponding to the possible detection of a lunar Askaryan pulse. For a valid trigger, we required an excursion beyond an adjustable threshold in the dedispersed data for a beam pointing at the Moon, in either polarisation; this excursion could occur either in the direct voltage samples, or in an equal number of values interpolated midway between them. Further, we required that the pulse

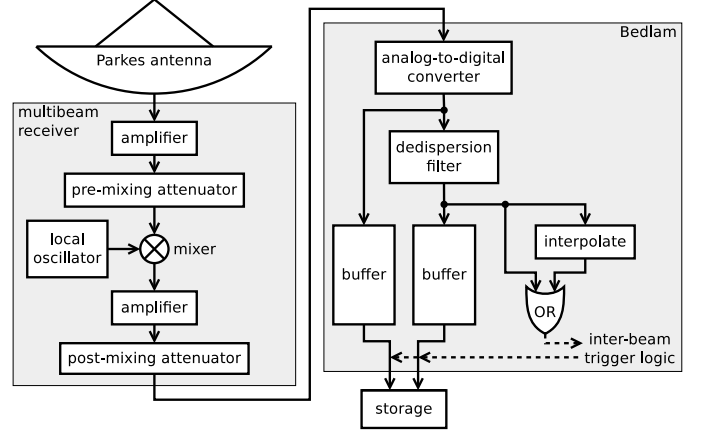


Figure 1: Signal path for a single telescope beam for this experiment, for one polarisation only. Left side is the signal path through the Parkes 21 cm multibeam receiver, which is standard except for the addition of an attenuator prior to the mixer (see Sec. 2.2). Right side shows the Bedlam backend built for this experiment. Anticoincidence trigger logic between the beams is shown in Fig. 2.

appear in only a single beam: simultaneous excursions in multiple beams are characteristic of transient RFI from artificial sources in the vicinity of the telescope, detected through the far sidelobes of the beam patterns. The trigger logic of this anticoincidence filter is illustrated in Fig. 2.

The raw trigger rate was highly variable and typically dominated by RFI, with a mean rate of over 200 Hz. The anticoincidence filter rejected the majority of these raw triggers, with the accepted trigger rate typically in the range 1–2 Hz, dominated by thermal noise. On an accepted trigger, the raw and dedispersed buffers for all beams were copied to permanent storage; these constitute a single frame of data. During this data transfer, the system was unable to respond to further triggers, resulting in a minor loss of effective observing time. The recorded data were subjected to retrospective processing to optimise the signal-to-noise ratio (Sec. 5), to remove minor instrumental effects (Sec. 6.1), and to exclude RFI pulses which escaped the real-time anticoincidence filter (Sec. 6.2).

2.1. Pointing strategy

The Parkes 21 cm multibeam receiver has a feed consisting of thirteen horn antennas in a hexagonal array, forming a corresponding pattern of beams on the sky. The feed array can be rotated within a limited range, which is typically used to maintain a constant parallactic angle but can be arbitrarily controlled with software developed for this experiment. Approximating the telescope antenna as a uniformly-illuminated 64 m aperture (see Sec. 3), each beam is a frequency-dependent Airy disk with a full-width half-maximum (FWHM) size of $12.3'$ at the centre of the band.

Our pointing strategy was dictated by limitations on the geometry of a detectable UHE particle interaction. The radiation from a particle cascade is directed forward as a hollow cone at the Cherenkov angle; combined with

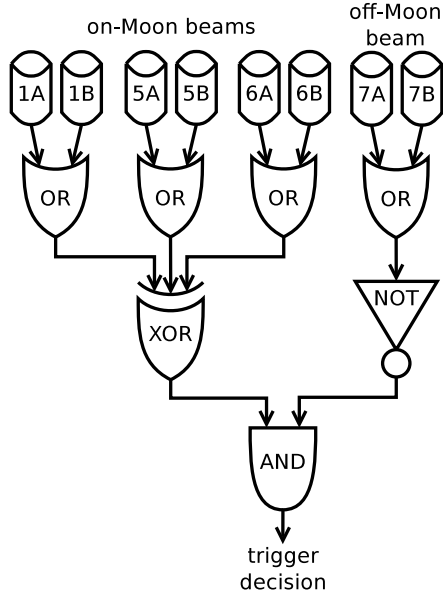


Figure 2: Real-time inter-beam trigger logic, for polarisations A and B of each of four beams. The choice of beams here is for pointing configuration A (see Fig. 3). In some configurations, two of the four beams were placed off-Moon, but the real-time trigger logic was not changed to reflect this.

internal reflection at the regolith-vacuum interface, this prevents the radiation from escaping for particles which are steeply down-going with respect to the lunar surface. Particles, even neutrinos, which are steeply up-going are strongly attenuated as they pass through the Moon. The only detectable particles are those which intersect the lunar surface at a shallow angle, with the refracted radio emission directed slightly above the surface. As viewed from the Earth, the Askaryan pulse comes from the Moon's apparent edge or limb, with the initial particle originating from a point typically $15\text{--}30^\circ$ in the corresponding direction from the Moon [22]. The component of the pulse which escapes the lunar surface has linear polarisation which is typically radial with respect to the Moon.

Combined with the consideration that the noise in our receiver is dominated by thermal radiation from the Moon, these limitations lead us to the following requirements.

1. As many beams as possible should be directed towards the limb of the Moon.
2. The Moon should occupy as little of each beam as possible, to minimise the thermal lunar noise in the receiver.
3. One linear polarisation of each beam should be oriented radially to the Moon, to match the polarisation of the expected signal.
4. To maximise our sensitivity to a potential UHE particle source, it should be close ($\lesssim 45^\circ$) to the Moon, and one beam should be directed at the point on the lunar limb closest to it.

Requirements 1 and 2 are contradictory, but it is possible to partially fulfill both of them by placing a beam

slightly off the limb of the Moon. The layout of the multi-beam receiver allows us to meet requirements 1–3 for two beams simultaneously, with both beams being slightly off-limb ($20'$ from the lunar centre, compared to a median lunar radius of $16'$) and having radial/tangential polarisation alignment (see Fig. 3). This allows a third beam to be placed on the Moon, although it is in a non-optimal half-limb position ($10'$ from the lunar centre) and has inferior sensitivity. The fourth beam supported by our backend hardware is placed away from the Moon, and serves to improve the efficacy of our anticoincidence filter. As it receives no significant contribution from the thermal radiation of the Moon, it is particularly sensitive and thus effective at detecting and identifying local RFI. Due to the importance of this role, we later changed the configuration to replace the half-limb beam with a second off-Moon beam, as shown for configurations A* and B* in Fig. 3.

Requirement 3 assumes that the Askaryan pulse maintains its polarisation angle from the Moon to the telescope, neglecting the Faraday rotation it will undergo as it passes through the ionosphere. The Faraday rotation angle for a thin-layer ionosphere, in radians, is

$$\beta = \frac{e^3}{8\pi^2 m_e^2 \epsilon_0 c} N_e B_{\parallel} \nu^{-2} \quad (1)$$

where ν is the radio frequency, e and m_e the charge and mass of the electron, ϵ_0 the vacuum permittivity, c the speed of light in vacuum, N_e the column density of free electrons in the atmosphere integrated along the line of sight, and B_{\parallel} the component along the line of sight of the geomagnetic field. For a typical geomagnetic field of $50 \mu\text{T}$, pessimistically assumed to be parallel to the line of sight, and $N_e \sim 10\text{--}20 \times 10^{16} \text{ electrons m}^{-2}$ (typical for our observations; see Sec. 4), the maximum rotation angle in our band is $5\text{--}10^\circ$, corresponding to a loss of signal strength of $\sim 1\%$ in the expected polarisation.

We applied requirement 4 to the potential UHE particle source Centaurus A, a nearby radio galaxy which is weakly correlated with the arrival directions of detected UHE cosmic rays [23]. This restricted our observations to a period of 3–5 days in each 27-day lunar cycle during which the Moon makes its closest approach to Centaurus A. It also required rotation of the feed array to keep one beam at the correct point on the lunar limb. Due to the feed rotation limit and the motion of the Moon across the sky, it is usually not possible to maintain the required pointing with the same set of beams for an entire observation. This occasionally made it necessary to switch to a different set of beams, for which the required pointing was within the available feed rotation range.

The required telescope pointing and feed rotation were determined at one-minute intervals with a combination of the SLALIB¹ and PYEPHEM² libraries. Each pointing was

¹<http://www.starlink.rl.ac.uk/docs/sun67.htx/sun67.html>

²<http://rhodesmill.org/pyephem/>

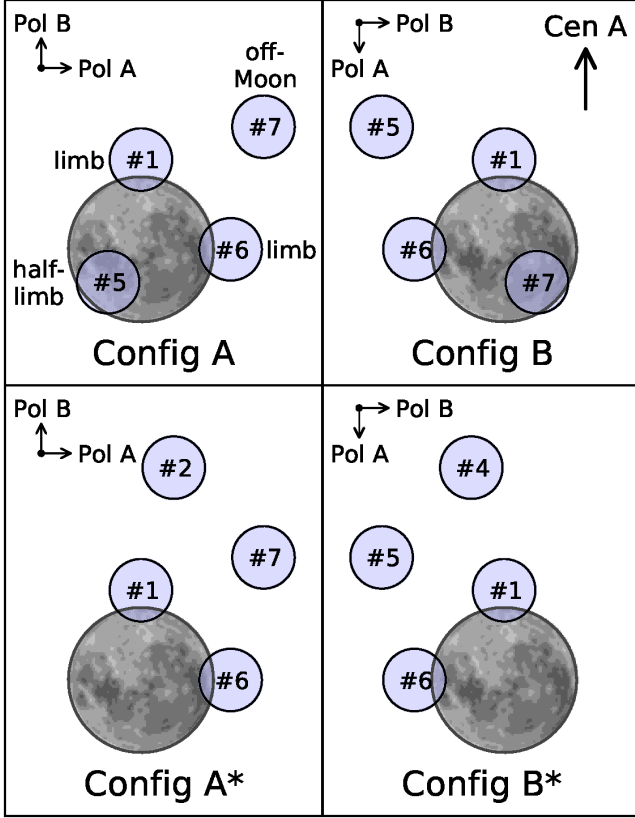


Figure 3: Pointing configurations employed for pulse detection in this experiment; an additional centre-Moon pointing (not shown) was used for calibration, as described in Sec. 3. Each configuration uses a (numbered) subset of the thirteen beams available with the Parkes 21 cm multibeam receiver. The limb and half-limb beam pointings (labelled, top left) are sensitive to radio pulses from the Moon, while the off-Moon beams are used for anticoincidence filtering. For each limb beam, one linear polarisation (“Pol A” or “Pol B”) is aligned radially to the Moon. In all configurations, the feed was rotated such that Centaurus A was towards the top of the figure, as shown; this requirement, and the rotation limit of the feed, required occasional changes between configurations A and B, or A* and B*. The separation from the lunar centre is $20'$ for the limb beams and $10'$ for the half-limb beam.

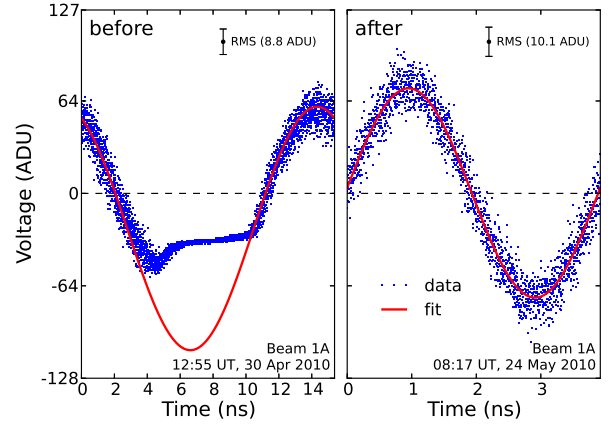


Figure 4: Single buffers of sampled values, recorded while transmitting a sine wave signal directly into the receiver, folded to the period of the signal. Left and right panels show data from before and after additional attenuation was inserted early in the signal path to restore linear behaviour at large voltage amplitudes. The solid line shows a sine-wave fit to the data near zero voltage, where the signal behaves linearly in both cases. The frequency of the signal is different in each case, but we do not expect this to significantly influence the measurement of the non-linearity effect.

a fixed position on the sky in equatorial coordinates; the motion of the Moon relative to the sky is the dominant source of pointing error, amounting to $\pm 0.25'$ within each interval, which is small compared to the beam size.

2.2. Receiver linearity

For an Askaryan pulse to be detectable by this experiment, it must have a peak amplitude substantially greater than the thermal noise from the Moon, exceeding the dynamic range for which radio telescopes are typically designed. It is therefore critical to ensure that the receiver, in translating the electric field of the received radio signal to the output signal voltage, maintains a linear response even at large amplitudes. Jaeger et al. [14] used only antennas of the EVLA in their experiment, despite the availability of additional unupgraded antennas of the VLA (Very Large Array), because they found that the latter failed to maintain the necessary linear response.

For our experiment, we tested the linearity of the response during our initial observations in April 2010, by transmitting a strong monotone signal from a small antenna within the 64 m telescope dish directly into the prime-focus receiver. Buffers of data were captured via our standard signal path and showed that the receiver was saturating, failing to properly amplify the sine-wave signal in the negative voltage direction (see Fig. 4). This problem was solved by inserting 20 dB of attenuation prior to the component in the signal path responsible for the non-linear behaviour, returning it to its range of linear operation, and adding an extra 20 dB of gain at a later stage (see Fig. 1). This hardware change was applied in May 2010 and for all subsequent observations.

2.3. System dispersion

The detection of a pulse coherent across a range of frequencies, as attempted in this experiment, requires that the components of the pulse at different frequencies arrive simultaneously, to a precision better than the inverse bandwidth. As conventional radio observations are unaffected by dispersion on this timescale, it is possible that some degree of such dispersion is inherent to the telescope system (not to be confused with ionospheric dispersion; see Sec. 4). Previous experiments have measured the impulse response of their receiver systems, but not to the precision required to demonstrate that this effect is negligible.

We tested the system dispersion with a network analyser, measuring the group delay through the receiver system at each frequency within the band, using the same test antenna as in Sec. 2.2. We found the group delay to have a root mean square (RMS) variation on a 1 MHz scale of 0.15 ns, which is a small fraction of the inverse bandwidth and hence of the expected pulse width. The resulting amplitude loss relative to an undispersed pulse, calculated by applying corresponding random delays to frequency components of a simulated pulse, is less than 0.01%.

3. Calibration

To establish the sensitivity of our experiment to an Askaryan pulse we must calibrate the response of our instrument, from the spectral electric field strength $\mathcal{E}(\nu)$ of the pulse through to the amplitude of the resulting peak in the signal $s(t)$ recorded in a buffer. We performed this calibration using the Moon, which is a strong source of thermal emission in our radio band and larger in angular size than our primary telescope beam, separately for each polarisation channel of each beam. This procedure was more complex than using an unresolved astronomical point source as a calibrator, but served the additional purpose of testing our assumptions about the beam shape and the distribution and polarisation of the lunar thermal emission.

The signal $s(t)$ in a recorded buffer of data from our experiment is measured in analog-to-digital units (ADU). For a buffer of length Δt , we find the power spectrum $P(\nu)$ from the signal's Fourier transform $S(\nu)$ as

$$P(\nu) = S(\nu)^2 / \Delta t \quad (2)$$

where the transform is normalised such that

$$\int d\nu P(\nu) = s_{\text{rms}}^2 \quad (3)$$

with s_{rms} being the RMS noise level. The power spectrum is related to the flux density $F(\nu)$ by

$$P(\nu) = B(\nu) F(\nu) \quad (4)$$

where the bandpass function $B(\nu)$ incorporates the unit conversion. The flux density contains contributions from

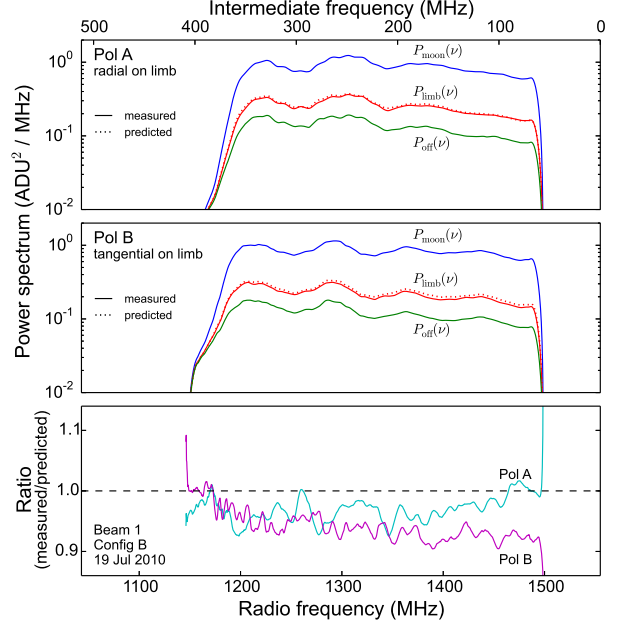


Figure 5: Measured power spectra for centre-Moon, off-Moon and limb pointings, as a function of both RF (bottom axis labels) and IF (top axis labels), for both polarisations (upper and centre panels) of a single beam. Variation between beams is similar to the variation between polarisations shown here. For the limb pointing, one polarisation is aligned radially to the Moon, and the other tangentially. Also shown (dotted) are the spectra predicted for the limb pointings. Finally, the lower panel shows the ratio between the measured spectra and these predictions.

the system noise of the receiver and the flux received from a strong source such as the Moon, the latter of which will vary depending on the position of the source in the beam. For calibration, we pointed each beam successively at the centre of the Moon, in the limb position shown in Fig. 3, and off the Moon. In each position, we recorded $\sim 1,000$ buffers over the course of a minute, comprising a few ms of data, and calculated averaged power spectra of $P_{\text{moon}}(\nu)$, $P_{\text{limb}}(\nu)$, and $P_{\text{off}}(\nu)$ respectively. These were subjected to a median filter with width 4 MHz to remove narrow-band interference, and a second-order Savitzky-Golay smoothing filter [24] with width 16 MHz to reduce the noise level while preserving features of this scale or larger.

The resulting spectra are shown in Fig. 5. We note that there is significant structure visible at a scale of ~ 50 MHz. As the thermal spectrum of the Moon contains no such structure, this must result from the bandpass function, which encompasses the effects of all components of the analog portion of the signal path (Fig. 1, left side). This bandpass structure is quite normal for such wideband analog systems and is quite stable in time. This frequency structure has no significant effect on the sensitivity of the experiment (see Sec. 5.1).

Among the analog components of the signal path are the post-mixing attenuators, which were adjusted after each change of pointing configuration to maintain equal

power levels in the swapped off- and on-Moon beams. We therefore performed this calibration routine once per pointing configuration per 3–5 day observing run, except for configuration A during 24–26 May 2010, which was calibrated based on an assumed equivalence between the received lunar thermal emission in each beam during this run and in the same configuration during the following month; the results are not notably inconsistent. On one occasion, we calibrated twice on the same day with the Moon at different elevations (56° vs. 79°) and measured $\sim 2\%$ variation in the signal power, which we ascribe to thermal radiation from the ground received through the far sidelobes of the telescope.

The contribution of the Moon to the received flux is found, for a single polarisation, with the Rayleigh-Jeans law:

$$F_{\text{moon}}(\nu) = \frac{k\nu^2}{c^2} \int d\Omega \mathcal{B}(\theta, \nu) T(\hat{\Omega}, \phi) \quad (5)$$

where k is the Boltzmann constant and $\mathcal{B}(\theta, \nu)$ is the beam power pattern, assumed to be radially symmetric (we take it to be an Airy disk) and normalised to $\mathcal{B}(0, \nu) = 1$; $B(\nu)$ is thus the bandpass at the centre of the beam, at $\theta = 0$. $T(\hat{\Omega}, \phi)$ is the brightness temperature of the Moon at position $\hat{\Omega}$. As the thermal emission from the Moon is partly linearly polarised, it must be expressed as

$$T(\hat{\Omega}, \phi) = (T_{\text{tot}}(\hat{\Omega}) - T_{\text{pol}}(\hat{\Omega})) + T_{\text{pol}}(\hat{\Omega}) \cos^2 \phi \quad (6)$$

in terms of the total and polarised components $T_{\text{tot}}(\hat{\Omega})$ and $T_{\text{pol}}(\hat{\Omega})$, with ϕ being the angle between the polarisation of the receiver and the polarisation of the emission. We take these brightness temperatures from Moffat [25], who fit a model of the lunar thermal emission to interferometric observations at 1,420 MHz. Their model incorporates the effects on the intensity and polarisation of the thermal emission from subsurface layers of the Moon as it is refracted through the rough lunar surface, as well as polar cooling.

The true equatorial temperature (i.e. the subsurface temperature before the emissivity of the surface is taken into account) in the model of Moffat [25] is 224 K. This value is approximately constant across a wide band, including the frequencies used in this experiment [26]. The model does not incorporate the $\lesssim 5$ K temperature difference between the lunar maria and highland regions, which leads to a $\pm 1\%$ uncertainty, nor the variation over a lunar day of about $\pm 3\%$ [26, and references therein]. The systematic error is somewhat larger, with an uncertainty of $\pm 8\%$ in the absolute calibration of Moffat [25], which is typical for radio observations.

With the received flux from the Moon calculated with Eq. 5, the received flux from the off-Moon sky assumed to be zero, and measurements of the power spectra $P_{\text{moon}}(\nu)$ and $P_{\text{off}}(\nu)$, the bandpass function $B(\nu)$ can be calculated. As a test, we use the calculated bandpass function to determine the expected power spectrum $P_{\text{limb}}(\nu)$ in the

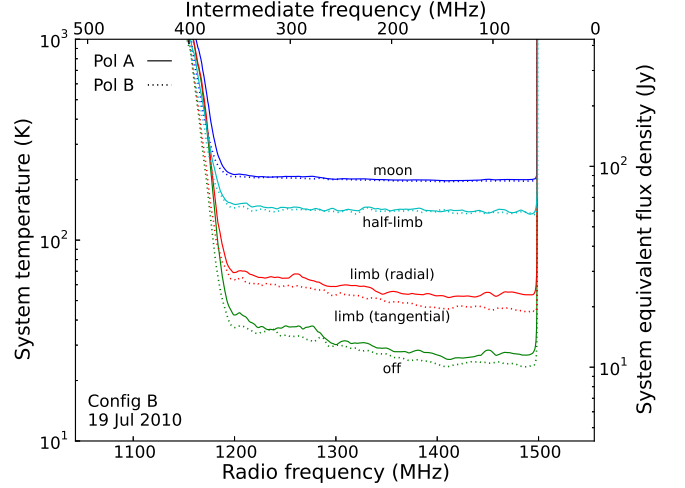


Figure 6: Calibrated spectral noise in terms of the system temperature (left) and the system equivalent flux density (right; for one polarisation only), for four different beam pointings with respect to the Moon. The limb and half-limb pointings are equivalent to beams #1 and #7 respectively in Fig. 3 (configuration B), and the remaining two cases are for pointing off the Moon (“off”) and at its centre (“moon”). The 1 Jy offset between the two polarisations for most pointings indicates a difference in the base receiver noise in the two channels; the increased 3 Jy offset for the limb pointing results from the alignment between polarisation A and the radial polarisation of the lunar thermal emission.

limb-pointing position, and compare it with the measured spectrum. A comparison of this type is shown in Fig. 5; in general, the results are consistent within the range of random error described above.

We conducted this calibration procedure for each pointing configuration of every observing run; Fig. 6 shows the calibrated spectral noise in one such case. The results in each case were internally consistent, but there were differences in the noise level between runs, primarily due to variation in the apparent size of the Moon and hence in the received thermal radiation. For simplicity in determining the overall sensitivity of this experiment to an Askaryan pulse, we adopt for each distinct beam position and polarisation a single value for the noise, averaging over this variation. We take the maximum measured variation as the error in the noise power, and hence in the sensitivity, which gives a random uncertainty of $\pm 9\%$, comparable to the $\pm 8\%$ systematic uncertainty. Both of these dominate over the error in the calibration procedure for each individual run.

The sensitivity to an Askaryan pulse is best expressed as a threshold in the spectral electric field strength immediately prior to being received by the antenna. As the electric field is proportional to the square root of the power, the uncertainty in this value will be half the uncertainty quoted above. To relate the electric field strength to the recorded signal $s(t)$, we use its relation [11] to the flux

density for a limited interval Δt :

$$F(\nu) = \frac{1}{Z_0 \Delta t} \mathcal{E}(\nu)^2 \quad (7)$$

where Z_0 is the impedance of free space. From Eqs. 2 and 4 we can then find that

$$S(\nu) = \sqrt{\frac{B(\nu)}{Z_0}} \mathcal{E}(\nu) \quad (8)$$

for a signal originating from the centre of the beam. Assuming the signal to comprise a coherent pulse with zero phase at all frequencies (i.e. $\arg(S(\nu)) = 0$; neglecting the effects described in Secs. 5.2 and 5.3), the peak amplitude is

$$s_{\text{peak}} = \int d\nu S(\nu). \quad (9)$$

Approximating the pulse spectrum $\mathcal{E}(\nu) = \mathcal{E}$ to be constant, and inserting an additional factor $\mathcal{B}(\theta, \nu)$ to allow for a signal originating away from the centre of the beam, Eqs. 8 and 9 then give us

$$\mathcal{E} = s_{\text{peak}} \left(\int d\nu \sqrt{\frac{B(\nu) \mathcal{B}(\theta, \nu)}{Z_0}} \right)^{-1} \quad (10)$$

as the threshold spectral electric field strength for the detection of such a pulse. Note that s_{peak} may also be represented as

$$s_{\text{peak}} = n_{\sigma} s_{\text{rms}} \quad (11)$$

in terms of the significance n_{σ} of the pulse relative to the RMS noise. The value of this threshold significance is found through analysis of the observational data in Sec. 6.

4. Ionospheric dispersion

A major consideration for this type of experiment is the dispersion of coherent pulses as they pass through the ionosphere, which causes a frequency-dependent delay in the pulse arrival time. For a wide bandwidth, the effect is to stretch the pulse out in time and to reduce its peak amplitude, particularly at low frequencies. Early experiments looked in multiple frequency bands for coincident pulses, separated by an offset corresponding to the dispersive delay [8, 10]. Others, operating at high frequencies with relatively narrow bandwidths, simply neglected the minor dispersive effects [9, 14]. Buitink et al. [13], whose experiment was the most severely affected by dispersion due to its lower frequency, recorded voltage data continuously over the course of their observations, and dedispersed it afterwards to restore the original pulse. James et al. [11] used an analog dedispersion filter developed by Roberts [27] which allowed the same correction to be performed in real time, but it was limited to a fixed dispersion characteristic, whereas the effect of the ionosphere varied over the course of the observations.

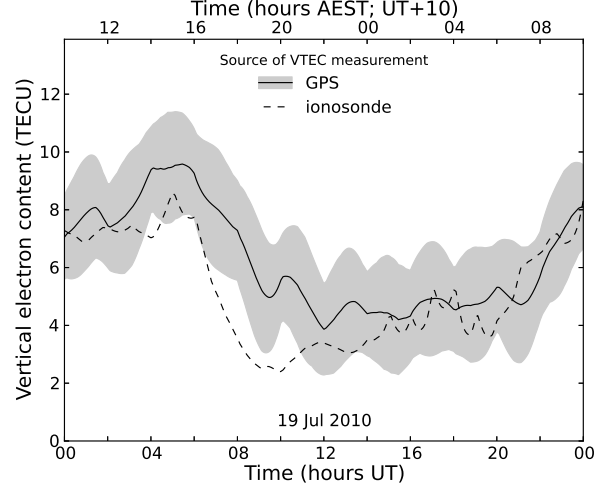


Figure 7: GPS- and ionosonde-derived VTEC over a single day at the Parkes telescope. The grey shading represents the uncertainty in the GPS measurement. No uncertainty bounds are available for the ionosonde data, but its deviation from the GPS values provides an indication of its precision.

For this experiment, we developed a digital dedispersion filter (see Sec. 5.2) which could be adjusted to compensate for different degrees of dispersion in real time. This allowed us to detect a dispersed pulse and trigger the storage of buffered data, so that more precise corrections could be applied in retrospective processing. To determine the correct setting for the filter, we required real-time information about the degree of dispersion, which is determined by the ionospheric total electron content (TEC). This is typically reported as a column number density in TECU units ($1 \text{ TECU} = 10^{16} \text{ electrons m}^{-2}$), either for a vertical column (vertical TEC, or VTEC) or a column along the line of sight (slant TEC, or STEC).

The free electron content of the ionosphere is controlled by ionisation caused by the high-energy component of solar radiation, and subsequent recombination with positive ions over time. As a result, it varies with latitude and season, and peaks several hours after local noon (see Fig. 7). It also varies with the 11-year solar magnetic activity cycle; our observations were conducted towards the end of the most recent solar minimum (see Fig. 8).

For our experiment, we used TEC values derived from two different sources. The first of these is the Global Positioning System (GPS), which measures the dispersion of radio signals between a network of ground stations and a constellation of satellites in medium Earth orbit. These values are the most accurate, but were only available retrospectively, after a delay of several weeks. The second source is ionosonde measurements, which use active radar to probe the reflectivity of the plasma in the lower ionosphere. These values are less accurate, and neglect conditions in the upper ionosphere, but were available in near real-time.

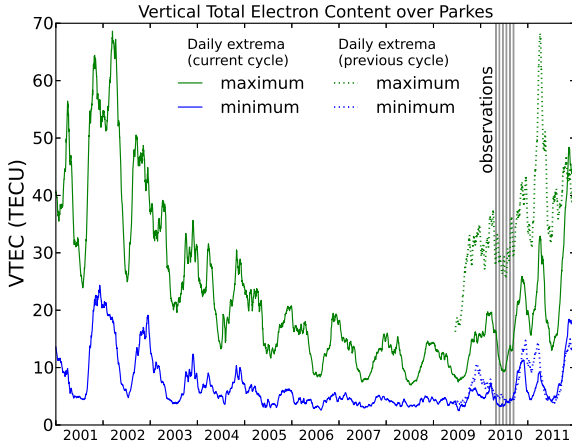


Figure 8: Electron content of the ionosphere over Parkes, derived from GPS data. Values shown are the maximum and minimum per 1-day cycle, smoothed with a width of 20 days, for the most recent 11-year solar cycle and the tail end of the previous one. At the time of our observations (vertical bars) the electron content was quite low, particularly compared to the equivalent point in the previous solar cycle.

4.1. GPS electron content measurements

Dual-frequency GPS measurements determine the STEC along the lines of sight between satellites and ground-based receivers to a precision of better than 0.1 TECU [28]. These STEC measurements at a sparse set of points, between each receiver and each visible satellite, are then interpolated to produce a global map of ionospheric VTEC by several groups; we used the maps generated by the Center for Orbit Determination in Europe (CODE). However, this procedure introduces substantial uncertainty, due both to the interpolation between sparse measurements and the assumption that the ionosphere can be represented with a single VTEC value at each coordinate, implying that it exists as a single thin layer. CODE also supplies maps of the estimated uncertainty in the VTEC measurement; during our observations, this was typically ± 2 TECU (see Fig. 7).

The VTEC maps were obtained from the Crustal Dynamics Data Information System (CDDIS)³ [29], in the IONEX (IONosphere map EXchange) format [30]. A single file in this format specifies the VTEC on a regular grid in geocentric latitude/longitude (i.e. on a spherical shell) at intervals over the course of a day. To find the STEC during our observations, we first determined the pierce point on this shell of the line of sight between the Parkes telescope and the Moon as shown in Fig. 9, and performed bilinear interpolation between spatially adjacent grid points to find the VTEC. We interpolated in time between the previous and next VTEC maps, rotating them in longitude to match the local time at the pierce point, using the

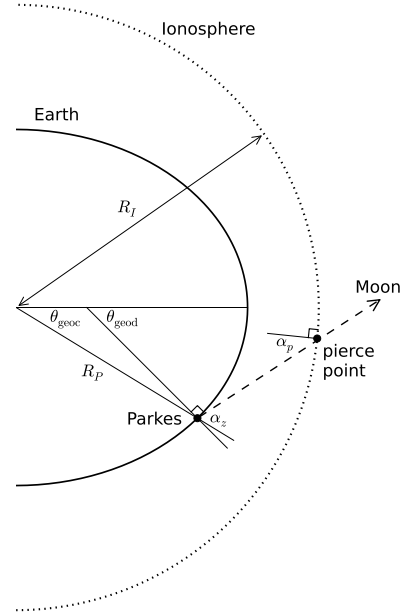


Figure 9: Determination of the ionospheric pierce point on the line of sight (dashed) from the Parkes telescope to the Moon, for the single-layer ionosphere model. Note the distinction between geodetic latitude θ_{geod} and geocentric latitude θ_{geoc} ; the zenith angle α_z is defined with respect to the geocentric vertical. The ratio between VTEC and STEC is determined by the angle of incidence α_p of the line of sight at the pierce point, as in Eq. 13. (Not to scale. The oblateness of the Earth has been exaggerated).

third and most precise interpolation method described by Schaer et al. [30]. Finally, we converted from VTEC to STEC using a slant factor based on the angle α_p between the line of sight and the normal to the ionosphere at the pierce point, given by

$$\frac{\text{STEC}}{\text{VTEC}} = \frac{1}{\cos \alpha_p} \quad (12)$$

$$= \left(1 - \left(\frac{R_P}{R_I} \sin \alpha_z \right)^2 \right)^{-1/2} \quad (13)$$

where α_z is the zenith angle relative to the geocentric vertical, and R_P and R_I are the radii from the centre of the Earth of the Parkes telescope and the model ionosphere shell respectively. The CODE maps assume a single-layer ionosphere at an altitude of 450 km relative to a mean Earth radius of 6,378 km, for a total $R_I = 6,828$ km.

The error introduced by the above procedure for determining the STEC from a CODE VTEC map should be substantially smaller than the map's reported uncertainty, so we display only the latter uncertainty elsewhere in this paper, scaled by the slant factor. This uncertainty was typically ± 2 TECU, with a maximum value of 3.8 TECU over the course of our observations. These values omit the dispersive effect of electron content at altitudes exceeding that of GPS satellites (20,200 km) but, applying a model of the high-altitude electron distribution [31], we find this to be $\lesssim 0.1$ TECU, which is negligible.

³[ftp://cddis.gsfc.nasa.gov/pub/gps/products/ionex/](http://cddis.gsfc.nasa.gov/pub/gps/products/ionex/)

4.2. Ionosonde electron content measurements

An ionosonde is a radar used to probe conditions in the ionosphere, typically operating in the frequency range 1–20 MHz. The maximum frequency at which it detects a reflection from the F2 layer of the ionosphere, referred to as foF2, is a measure of this layer’s plasma frequency and hence of its electron density. The Ionospheric Prediction Service (IPS), a branch of the Australian Bureau of Meteorology, makes available IONEX-format maps of VTEC based on foF2 measurements from a network of ionosondes⁴.

As foF2 directly measures only the peak electron density in the F2 layer, rather than the integrated column electron density, ionosonde data are less accurate than GPS-derived TEC values. Fig. 7 shows the variation between data from these two sources. The advantage of the ionosonde VTEC maps, however, is that they are available in near real-time, with a delay of only 1–2 hours. This allowed us to use them to calibrate our dedispersion filter during our observations. We derived STEC values from VTEC maps as for GPS data in Sec. 4.1, and estimated the current value as

$$\text{STEC}_{\text{est}}(t) = \text{STEC}(t - \Delta t) \times \frac{\text{STEC}(t - 24 \text{ hrs})}{\text{STEC}(t - 24 \text{ hrs} - \Delta t)}, \quad (14)$$

scaling the most recent available value (at time $t - \Delta t$) by the change in STEC during the equivalent period on the previous day.

As no error estimate is available for the ionosonde-derived TEC values, we take as the error the difference between them and the more accurate values derived from GPS measurements, which was typically ± 4 TECU in STEC. During our observations, we set the dedispersion filter to correspond to the multiple of 3.02 TECU (corresponding to 1 ns of dispersion across the RF band 1.2–1.5 GHz; see Sec. 5.2) closest to the ionosonde-derived STEC value. The resulting variation between the filter setting and the GPS-derived STEC is shown in Fig. 10 for a typical observing run; the maximum value of this error over the course of our observations was 9.0 TECU. We describe how we dedispersed the signal in Sec. 5.2, and the consequences of the STEC error and uncertainty for the sensitivity of our experiment in Sec. 5.5.

5. Signal optimisation

The signal-processing problem for an experiment of this type is to detect an Askaryan pulse against a Gaussian noise background dominated by thermal radiation from the Moon and internal noise in the telescope receiver (apart from RFI, which we address in Sec. 6.2). Previous experiments have either converted the signal to the power domain or searched directly in the voltage domain,

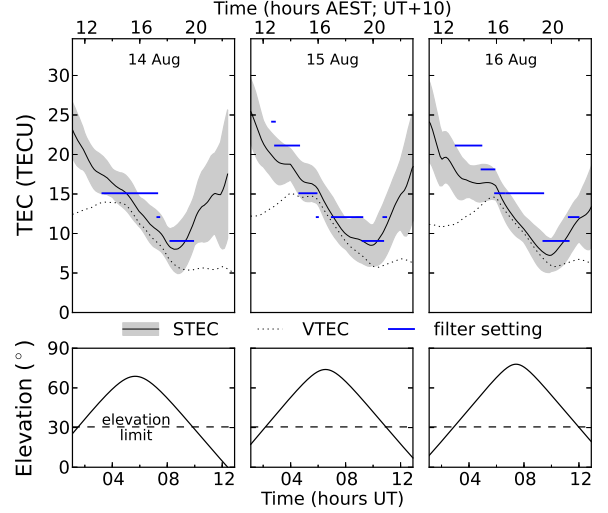


Figure 10: Comparison between ionospheric TEC and the STEC for which our dedispersion filter was set during our three-day observing run in August 2010. The upper plots show the GPS-derived STEC along the line of sight to the Moon (solid; uncertainty range shaded) and VTEC (dotted), as well as the filter settings (horizontal bars) with which we attempted to match the STEC. Ultimately, the filter settings were sufficiently closely matched to the STEC that there was no detriment to the retrospective sensitivity of the experiment (see Sec. 6.3). The lower plots show the elevation of the Moon relative to the elevation limit of the Parkes radio telescope; the increased STEC relative to the VTEC at the start and end of each observation is due to the Moon being close to the horizon.

using a range of filters to improve the signal-to-noise ratio. Here we use the voltage-domain approach, using the theoretically optimum[e.g. 32, Ch. 42] filtering strategy: the maximum signal-to-noise ratio is achieved by applying a pre-whitening filter with a bandpass that gives the noise a flat (white) spectrum, followed by a filter matched to the expected shape of the pulse. In terms of their respective transfer functions $W(\nu)$ and $M(\nu)$, we can represent these as a single combined filter with transfer function

$$O(\nu) = W(\nu)M(\nu). \quad (15)$$

We divide our discussion of the optimising filter $O(\nu)$ into three parts. In Sec. 5.1, we describe the derivation of its amplitude $|O(\nu)|$, which represents the optimum bandpass of the filter. In Sec. 5.2, we consider the derivative of the phase $\arg(O(\nu))$ across the band: the first-order term or linear phase slope corresponds to an absolute delay in the arrival time of a pulse, which does not affect its detectability, but the second- and higher-order terms describe the dispersion resulting from the passage of the pulse through the ionosphere (see Sec. 4), which is reversed by an appropriate filter. In Sec. 5.3, we consider the base phase $\arg(O(0))$ of the filter, which relates to the inherent shape of the pulse.

The optimising filter described above maximises the pulse amplitude relative to the background noise for an analog signal. In the digital domain, in which the filter is implemented, the signal is represented by a series of

⁴[ftp://ftp.ips.gov.au/data/Satellite/](http://ftp.ips.gov.au/data/Satellite/)

discrete samples and must therefore be interpolated to restore the full pulse amplitude, as described in Sec. 5.4. In our experiment, both the filter and this interpolation were applied in an approximate form in real time, and more precisely in retrospective processing. In Sec. 5.5 we discuss the signal loss in each case from the combination of the above effects.

5.1. Bandpass optimisation

Prior to filtering, the signal has the noise power spectrum $P(\nu)$, from Eq. 2. To give the signal a flat noise spectrum, the pre-whitening filter must have a bandpass

$$|W(\nu)|^2 \propto 1/P(\nu). \quad (16)$$

The purpose of this filter is to place a greater weight on the signal at frequencies with lower noise power, at which the signal-to-noise ratio is higher and the system is therefore more sensitive.

The pre-whitening filter should optimally be followed by a matched filter with transfer function $M(\nu)$, with its bandpass $|M(\nu)|^2$ matched to the spectrum of the expected pulse. If the pulse initially has an electric field spectrum $\mathcal{E}(\nu)$, and hence a power spectrum $\mathcal{E}(\nu)^2$, it will be modified first by the bandpass $B(\nu)$ of the receiver, and then the bandpass $|W(\nu)|^2$ of the pre-whitening filter, so

$$|M(\nu)|^2 \propto |W(\nu)|^2 B(\nu) \mathcal{E}(\nu)^2 \quad (17)$$

$$\propto \frac{\mathcal{E}(\nu)^2 B(\nu)}{P(\nu)}. \quad (18)$$

The bandpass of the combined optimising filter is then

$$|O(\nu)|^2 \propto (|W(\nu)||M(\nu)|)^2 \quad (19)$$

$$\propto \frac{\mathcal{E}(\nu)^2 B(\nu)}{P(\nu)^2} \quad (20)$$

which gives us the amplitude of its transfer function $O(\nu)$.

We applied this correction in retrospective processing separately for each beam, using the bandpass $B(\nu)$ derived in Sec. 3, and taking $\mathcal{E}(\nu)$ to be the electric field spectrum for a cascade observed in the most optimistic case, from the Cherenkov angle, in order to minimise the threshold detectable particle energy [33]. Typical improvements in the signal-to-noise ratio were 1–2%, and not sensitive to our assumption regarding the original pulse spectrum, resulting instead from corrections to minor variations in the bandpass and noise spectrum. This improvement is insignificant compared to the calibration uncertainty (see Sec. 3), indicating that these variations (see Fig. 5) have a negligible effect on the sensitivity, and we neglect this improvement in subsequent calculations.

The effect of bandpass optimisation will be greater for future experiments with larger fractional bandwidths. The significance of the assumption regarding the pulse spectrum will also increase, and it may become necessary to choose between optimising for maximum sensitivity to less

energetic particles, as we have here; or optimising for maximum aperture to more energetic particles, by biasing the filter towards low frequencies. This latter possibility is a more general case of a result found by James and Protheroe [5] under the assumption of a flat bandpass: that the aperture of a lunar radio experiment is improved under some circumstances by simply excluding the high-frequency end of its observing band.

5.2. Dedispersion

For radio waves significantly above the plasma frequency ($\nu_p \lesssim 10$ MHz in the ionosphere), the dispersive time delay relative to a signal traveling at the speed of light is

$$\Delta t = \frac{e^2}{8\pi^2 \epsilon_0 m_e c} N_e \nu^{-2} \quad (21)$$

$$= 1.34 \times 10^{-7} \left(\frac{N_e}{\text{m}^{-2}} \right) \left(\frac{\nu}{\text{Hz}} \right)^{-2} \text{ s}. \quad (22)$$

Note the similarity of Eq. 21 to the expression for Faraday rotation, another ionospheric effect, in Eq. 1: both are linearly dependent on N_e , the ionospheric STEC. The delay Δt is equivalent to multiplying the signal, in the frequency domain, by a phase factor $e^{i\phi}$ where

$$\phi = -2\pi \int_{\infty}^{\nu} d\nu' \Delta t \quad (23)$$

is the integral from an infinite frequency at which $\Delta t = 0$, though the choice of this limit affects only the base phase, not the dispersion. The delay can be reversed, and the signal dedispersed, by applying the conjugate phase factor $e^{-i\phi}$. If the signal is shifted in frequency prior to dedispersion, as in this experiment, the frequency ν in Eq. 21 represents the original radio frequency at which dispersion occurs, while the integral in Eq. 23 is instead performed over the intermediate frequency at which the dedispersion filter is applied.

5.2.1. Filter design

To perform this correction in full, a filter should transform the signal to the frequency domain, multiply by the appropriate phase factor, and transform the result back to the time domain, and we did this in retrospective processing. However, this was computationally impractical for us to achieve in the several μs available in real time, to detect a pulse while it was still present in the buffered data. The filter shown in Fig. 1 implements a convolution of the digital sampled data, in the time domain, with 64 arbitrary coefficients which describe its impulse response; in engineering parlance, it is a 64-tap finite impulse response (FIR) filter, with a delay between taps equal to the sampling interval of 0.98 ns. For this filter to act as an adjustable dedispersion filter, we needed to precompute and store an appropriate set of coefficients for each amount of dedispersion it was to apply.

To determine the optimum filter coefficients, we first calculated the required phase factor at a set of discrete

frequencies separated by $\Delta\nu = 0.5$ MHz, and applied an inverse discrete Fourier transform to find the corresponding impulse response. The length of this impulse response is $1/\Delta\nu = 2 \mu\text{s}$, or $\sim 2,000$ taps. As our filter only supports 64 taps, we truncated the impulse response to this length, selecting the set of 64 consecutive coefficients which maximised the retained power (i.e. the sum of the squares of the selected values), which we term the filter efficiency. The filter efficiency decreased significantly when calculating coefficients for a greater degree of dedispersion ($\gtrsim 100$ TECU), for which the required impulse response extends over a longer time and exceeds the 64-tap limit.

Using the above procedure, we generated sets of filter coefficients corresponding to dispersive delays of multiples of 1 ns across the radio frequency range 1.2–1.5 GHz (equivalent to multiples of 3.02 TECU in STEC), allowing us to recalibrate the filter during the experiment by loading new sets of coefficients. Fig. 11 shows the impulse response and bandpass of two such sets of coefficients, as well as the accuracy of the resulting dedispersion. The filter performance declined for dispersive delays greater than ~ 30 ns, due to loss of efficiency from truncation of the impulse response. However, ionospheric TEC during our observations was significantly lower than expected from the previous solar activity cycle (see Fig. 8), so the maximum filter setting used in our experiment was only 8 ns, for which the filter efficiency was $> 99.9\%$. The signal loss in dedispersion is instead dominated by uncertainty in the ionospheric TEC, as discussed in Sec. 5.5.

5.3. Pulse phase

Previous experiments of this type, when searching for a pulse in the voltage signal, have searched for an excursion in the voltage beyond some threshold magnitude. This implicitly makes the same assumption as in Eq. 9: that the signal has a spectrum $S(\nu)$ with zero phase, so in the time domain it combines coherently in a single direction. Experimental tests of the Askaryan effect show that this is not the case: the phase is closer to -90° , and it manifests in the time domain as a bipolar pulse [34], consistent with more recent simulations [35]. In theoretical terms, the two poles of the pulse can be understood as originating from the beginning and end of the particle cascade; the order in which they are observed depends on whether the observer is on the inside or outside of the Cherenkov cone. More precisely, one pole represents the increase in excess negative charge of the cascade, and the other represents its decrease: the pulse profile is the derivative of the charge excess along the length of the cascade [36]. The practical effect is that the power of the pulse, rather than being concentrated into a single unidirectional excursion of the electric field, is split between the two poles, and its maximum amplitude is decreased by a factor $\sim \sqrt{2}$.

A filter matched to the expected pulse would apply a phase of 90° (i.e. a Hilbert transform). However, in our experiment the phase is randomised when the RF signal is downconverted to IF by mixing it with a local oscillator

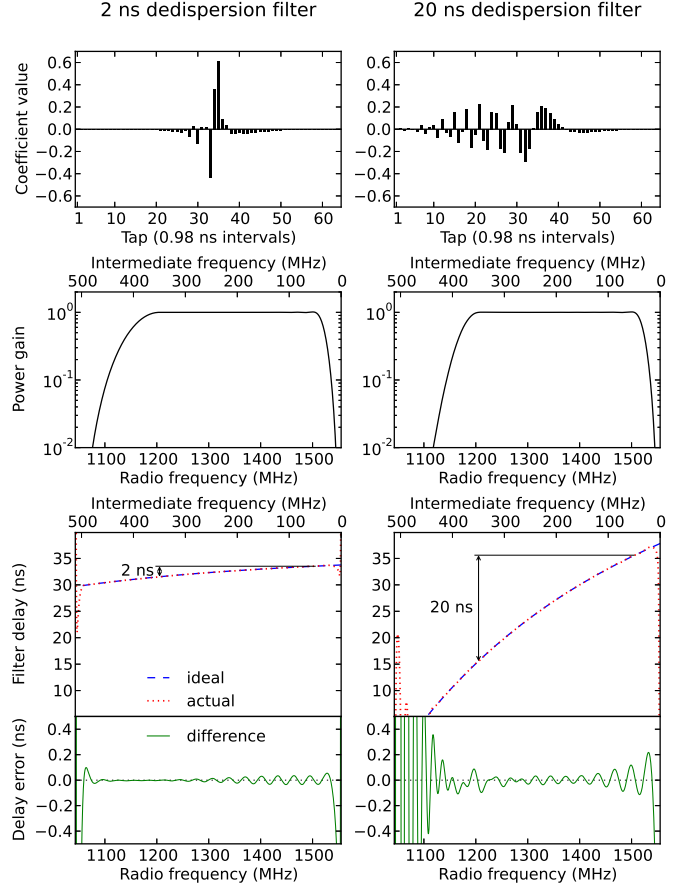


Figure 11: Performance characteristics of our real-time dedispersion filter when set to compensate for dispersion of 2 ns (left) and 20 ns (right) across the 1.2–1.5 GHz radio frequency range of our receiver. In the impulse response (top), there is a visible separation of the rapidly-oscillating high-frequency components (tap $\lesssim 32$) and the low-frequency components (tap $\gtrsim 32$): this corresponds to a greater group delay at high radio frequencies, counteracting the greater delay at low radio frequencies from ionospheric dispersion. The ideal (dashed) and actual (dotted) values of the group delay (bottom) correspond closely within the useful band; the deviation between the two (solid) is greater for the 20 ns filter than for the 2 ns filter, but is in either case typically < 0.1 ns. The absolute value of the delay is arbitrary, and corresponds roughly to half the length of the filter. The periodic oscillation in the delay error plot is similar to the behaviour of the analog filter developed by Roberts [27].

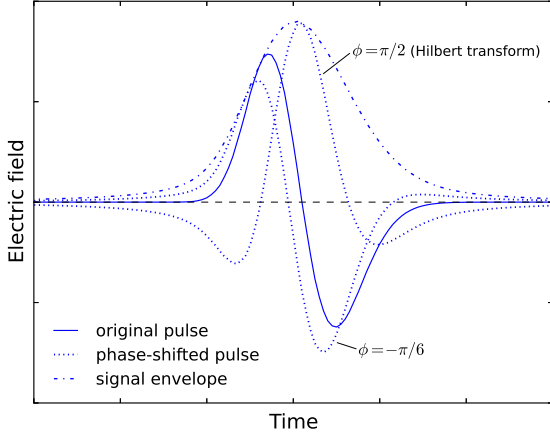


Figure 12: An idealised Askaryan pulse (solid), represented here as the derivative of a Gaiser-Hillas function [38]. Applying an arbitrary phase ϕ results in a range of other possible pulse shapes, two of which are shown (dotted). The signal envelope (dash-dotted) reconstructs the maximum possible amplitude of the pulse.

of unknown phase [21], and the resulting pulse will have one of a range of possible profiles (see Fig. 12). Instead, in retrospective processing of the data, we reconstruct the maximum possible amplitude of the pulse using the signal envelope as defined by Longuet-Higgins [37]. A side effect of this procedure is to increase the noise level, as the signal envelope has a Rayleigh rather than a Gaussian distribution of amplitude, which must be taken into account when determining the significance of a pulse.

5.4. Interpolation

The digital representation of a signal consists of a series of values corresponding to the voltage of the original analog signal at evenly-spaced sampling times. James et al. [11] note that, for a finite sampling rate, there will be a random offset between the peak of a pulse and the closest sampling time, which reduces the sensitivity of an experiment of this type, as the sampled digital data will not record the full magnitude of the pulse. They calculate the loss of sensitivity this effect causes for their experiment, finding it to be $\sim 0\text{--}30\%$, depending on the actual offset between the peak of the pulse and the closest sampling time, and on the coherency of the pulse; other similar experiments have neglected this effect entirely.

In this experiment, we correct for this effect by interpolating (up-sampling) the digital data, reconstructing the values of the analog signal at the times which were not originally sampled. Perfect interpolation requires a convolution of the signal with a sinc function to reconstruct the signal at every intermediate point. The real-time hardware used in our experiment convolves the signal with a digital approximation to a sinc function, interpolating a single value between every pair of samples. This doubles the effective sampling rate, halving the expected offset be-

tween the peak and the closest sampling time. Approximating the pulse as a sinc function, which has quadratic behaviour around its peak, the effect of this is to decrease the loss of sensitivity relative to the case of perfect interpolation by a factor of four. For retrospective processing, we performed 32-fold interpolation in software, with the remaining loss of sensitivity being negligible (reduced by a factor 32^2).

5.5. Signal loss

In our experiment, in real time, we applied a filter which dedisperses the signal (see Sec. 5.2.1), manually adjusting the filter setting based on ionosonde-derived STEC values (see Sec. 4.2), and performed two-fold interpolation to double the effective sampling rate. In retrospective processing, we applied full bandpass optimisation and more precise dedispersion based on GPS-derived STEC values (see Sec. 4.1), formed the signal envelope to search through possible phases of the signal, and performed effectively complete (32-fold) interpolation. To determine the potential loss of signal strength in each case, we simulated the progress of test pulses with different degrees of dispersion through the corresponding signal path, with the results shown in Fig. 13. The maximum loss of signal strength corresponds to the worst-case values for the signal phase and sampling times. Due to its minimal effect, we omitted bandpass optimisation from our simulation. Noise was omitted from our simulation: since the noise has no inherent phase or dispersion, it is equally valid to consider the effects of noise before or after processing.

The effects described in Secs. 5.2–5.4 are not clearly separable in their influence on the loss of signal strength, but we can give some indication of their relative significance by describing the magnitude of each of them individually, with the other effects completely removed. If the sampled digital data in this experiment had been subject to no processing other than a direct threshold test, the maximum signal loss would have been 17.9% due to the unknown pulse phase, 21.6% due to the finite sampling rate, or 15.0% due to dispersion for the maximum STEC of 23.5 TECU during the entire observations; the maximum total signal loss would have been 41.9% from these three effects combined (Fig. 13, “uncorrected”). With the actual processing implemented in real time, the loss due to the unknown pulse phase remains 17.9%, but two-fold interpolation reduces the sampling loss to 5.6%, and the dedispersion filter reduced the dispersive loss to 2.3%, calculated based on the maximum difference of 9.0 TECU between the filter setting and the STEC values retrospectively derived from GPS data; the total loss from these effects is 23.1% (Fig. 13, “real-time”). In retrospective processing of recorded data, the potential losses due to pulse phase and sampling are effectively eliminated by forming and interpolating the signal envelope, so the only remaining loss is 0.4% due to dispersion, calculated based on the maximum uncertainty of 3.8 TECU in the GPS-derived STEC values used for retroactive dedispersion (Fig. 13,

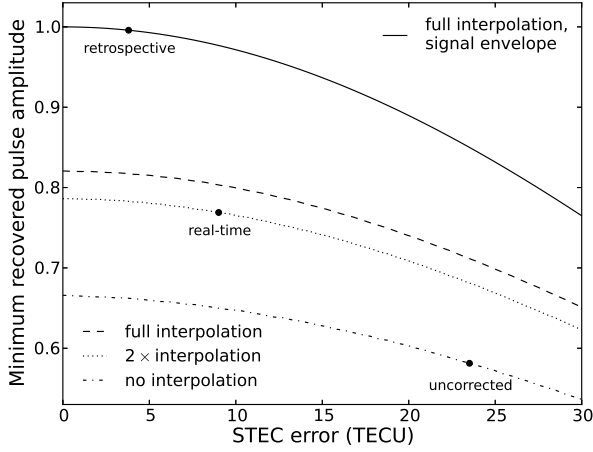


Figure 13: Maximum loss of sensitivity from suboptimal signal-processing in our experiment. In all cases, the sensitivity decreases if there is an error between the ionospheric STEC assumed by the dedispersion filter and the actual dispersion of the pulse. The solid line shows the recovered pulse amplitude after the full retrospective processing applied to recorded data, in which we perform effectively complete interpolation and form the signal envelope. The point labelled “retrospective” marks the value based on the maximum uncertainty of the GPS-derived STEC used in this processing. The dashed line shows the minimum recovered pulse amplitude if we had neglected to form the signal envelope, for the worst-case signal phase. The dotted line shows the minimum recovered pulse amplitude in real-time processing, in which only two-fold interpolation was performed. The point labelled “real-time” marks the value for the worst-case error in the dedispersion filter setting, equal to the maximum difference between the filter setting and the GPS-derived STEC. The dash-dotted line shows the minimum recovered pulse amplitude with no processing apart from dedispersion, for the worst-case signal phase and sampling times. The point labelled “uncorrected” marks the value with the absence of any dedispersion, for the maximum STEC during our observations.

“retrospective”). The relation of these signal losses to the final sensitivity of this experiment is discussed in Sec. 6.3.

6. Observations and analysis

Our observations were conducted from April to September 2010 in six runs of 3–5 days each month, using the observing strategy outlined in Sec. 2 and the calibration procedure described in Sec. 3. From 259.5 hours of scheduled time with the Parkes radio telescope, we obtained 148.7 hours of calibrated on-source data. Details of our observation schedule are in Table 1. The length of the buffer recorded on a successful trigger was set to $4\ \mu\text{s}$, except for the May observing run, during which it was set to $2\ \mu\text{s}$.

Attenuation was set individually for each beam and polarisation to maintain an RMS signal power of ~ 10 – 11 ADU, despite variation in system temperature between on- and off-Moon beams. This allowed peaks of up to $\sim 12\sigma$ to be recorded without exceeding the 8-bit digitisation range of -128 to $+127$ ADU. The attenuation settings were not adjusted to compensate for minor fluctuations in the instrumental gain, but were adjusted each time the pointing configuration was changed, moving beams between on- and off-Moon positions.

To permit the real-time anticoincidence filter to operate between the beams of the multibeam receiver, we calibrated their relative timing in each observing session using short-duration RFI pulses which were detected by all beams. The precision of this calibration was determined by the timescale of the RFI pulses: typically 5–10 ns, which is small compared to the 200 ns exclusion window applied around each trigger by the anticoincidence filter. The trigger threshold for the off-Moon beam was maintained at $\sim 4.6\sigma$, at which level the trigger rate was dominated by thermal noise, typically being ~ 10 kHz for both polarisations combined, although this varied significantly with small changes in the noise level. This trigger rate leads to $\sim 0.2\%$ of the observing time being excluded by the real-time anticoincidence filter (see Sec. 6.3.1). When a second off-Moon beam was employed (in configurations A* and B*; see Fig. 3), it was still treated as an on-Moon beam by the trigger logic (see Fig. 2), so it was set to a high threshold to prevent it from triggering and used for the exclusion of RFI only in retrospective processing (Sec. 6.2).

The threshold for triggers in the on-Moon beams, which cause events to be recorded, was maintained at $\sim 6.4\sigma$. Pulses exceeding this threshold were typically dominated by impulsive RFI, with a highly variable raw trigger rate which averaged 251 Hz but occasionally exceeded 100 kHz, based on logs of the raw trigger rates during each second. The majority of these raw triggers were excluded by the anticoincidence filter as shown in Fig. 14, and the events accepted by the filter were dominated by thermal noise, except for some short periods (total of several hours) during the September observing run with exceptionally high

Table 1: Details of our observations in 2010, including pointing configurations (see Fig. 3) and numbers of triggers in on-Moon beams.

Date	Time (AEST)		Duration ^a (hrs)	Configs	Real-time triggers		Triggers over 8σ	
	start	end			raw	accepted	accepted	passed cuts
27 Apr ^b	23:32	02:36	3.1	B	179 017	37 563 ^c	12	4
28 Apr ^b	20:23	04:27	7.1	B	292 925	18 952 ^c	9	0
29 Apr ^b	21:30	05:27	7.5	B	16 075 833	37 538	112	5
30 Apr ^b	22:00	06:33	7.1	B	16 288 167	64 257	78	0
24 May	18:51	01:15	5.8	A, B	17 042 475 ^d	35 898	34	6
25 May	18:12	02:09	7.5	A, B	16 483 399	59 140	399	4
26 May	18:26	03:19	8.7	A, B	452 461	65 551	22	4
27 May	19:09	02:50	7.5	B	2 102 839	60 902	69	3
20 Jun	15:03	23:07	7.5	A, B	1 014 873	25 962	148	5
21 Jun	15:35	00:10	8.3	A, B	723 232	32 432	164	5
22 Jun	17:21	00:12	6.3	A, B	1 354 033	13 447	68	7
23 Jun	17:09	01:51	8.4	B	930 107	35 207	55	4
17 Jul	18:10	20:59	2.2	B	348 323	14 434	11	2
18 Jul	13:49	22:04	7.7	A, B	2 370 062	44 735	101	4
19 Jul	15:01	23:05	7.7	A, B	20 931 824	57 788	60	2
20 Jul	14:59	00:06	8.8	A, B	8 812 410	57 052	173	3
21 Jul	16:02	01:02	8.9	B	13 716 161	62 076	59	2
14 Aug	13:17	19:53	5.8	A, B	5 302 640	32 248	163	3
15 Aug	12:39	20:57	7.2	A, B	197 411 ^d	25 872	36	2
16 Aug	13:02	21:57	8.7	A, B	4 724 087	36 600	92	6
10 Sep	09:53	17:32	2.4	A*, B*	1 437 656	9 994	136	0
11 Sep	10:11	18:41	8.3	A*, B*	9 323 492	38 235	798	3
12 Sep	10:49	19:45	8.7	A*, B*	3 206 896	38 709	650	4
13 Sep	11:43	20:48	4.0	B*	1 094 569	15 178	140	1
14 Sep	12:32	21:46	8.4	B*	11 571 661	33 108	927	3
Total ^b			148.7	A, A*, B, B*	123 140 611	794 568	4 305	73

^a Due to time lost to configuration changes etc., the duration of active observations is typically slightly less than the time between the start and end of an observing run.

^b Total does not include observations from April, which were excluded due to several technical problems (see text).

^c Real-time anticoincidence logic was not fully active for these runs.

^d Raw trigger rates were logged for only part of these observations, and have been extrapolated to fill the gaps.

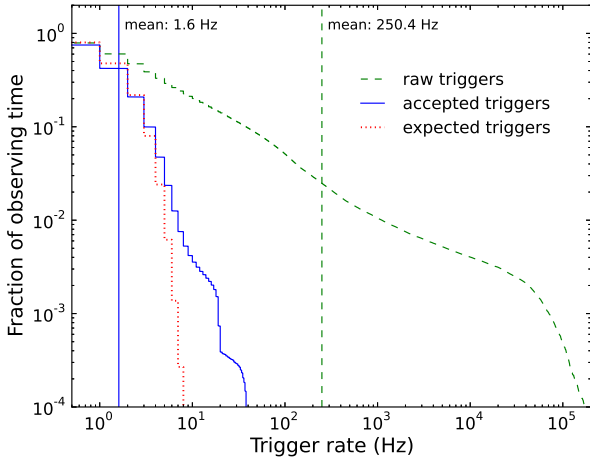


Figure 14: Cumulative histograms of the collective trigger rate of all on-Moon beams during our observations, both before the real-time anticoincidence filter (raw triggers; dashed) and after (accepted triggers; solid). The expected distribution of trigger rates is also shown (dotted), based on Poisson statistics within each 1 s interval and with the same mean rate as the accepted triggers. The raw trigger rate is typically elevated due to RFI, but most of this excess is eliminated by the real-time anticoincidence filter, with the rate of accepted triggers conforming much more closely to the expected distribution.

RFI activity, which were excluded at the time of the observations and not included in the analysis. The remaining thermal noise events occurred at a collective rate of 1–2 Hz for all on-Moon beams, and were recorded for further analysis.

Apart from the signal optimisation described in Sec. 5, we applied additional processing to recorded events to correct for minor instrumental effects, described in Sec. 6.1. High-significance events ($\gtrsim 8\sigma$) were dominated by RFI, so we applied various anti-RFI cuts to remove these, as described in Sec. 6.2. The remaining high-significance events occur at a low, constant rate which is consistent with the expected thermal noise. The amplitude of the most significant event to pass all cuts defines the threshold sensitivity of this experiment, which we discuss in Sec. 6.3.

6.1. Processing

Our data were affected by several minor bugs in the digital signal-processing firmware. One occasionally caused a single sample value to be repeated, overwriting the remainder of a buffer; another caused pairs of adjacent samples to be swapped, rendering the output of the dedispersion filter incorrect. Both of these affected only the initial observing run in April 2010, which was excluded from this analysis. A persistent bug caused spurious values to appear at the start and end of each buffer, but the samples which meet the trigger condition occur in the centre of the buffer and are therefore unaffected, so we simply truncated the first and last eight samples in each case.

As the dedispersion filter is a novel system developed for this experiment, near the limits of the capabilities of

the available digital signal-processing hardware, we tested the recorded data to ensure that it had performed as intended. We were able to do this because each frame of recorded data, corresponding to a single event which passed the real-time anticoincidence filter, contains copies of the signal both before and after it was acted on by the filter, allowing us to apply a software reimplement of the filter to the former copy to reproduce the latter. Doing this, we found that the filter produced inconsistent results under certain circumstances.

- In the April 2010 observing run, due to the sample-swapping bug described above, the filter acted on an out-of-order representation of the raw data, resulting in incorrect dedispersion. This bug, however, already motivated the exclusion of this run from this analysis.
- In the March 2010 observing run, a small fraction ($\sim 2\%$) of recorded frames showed a minor mismatch between the output of the real-time dedispersion filter and its software implementation acting on the raw data, with occasional samples differing by 1–2 ADU ($\sim 0.1\text{--}0.2\sigma$). The cause of this bug is unknown, but we consider it insignificant and ignore it.
- On each occasion when the dedispersion filter setting was changed, a single frame was not processed correctly under either the old or the new setting. These occasions were extremely infrequent, and excluded from further analysis.
- As both the raw and filtered data were recorded as 8-bit integers, the filter was unable to function correctly when dedispersion would increase the amplitude of a peak beyond the range -128 to $+127$ ADU ($\pm \sim 12\sigma$). These events were still recorded, and the raw data is still valid.

Partly because of this last issue, further analysis was performed only on the raw (rather than dedispersed) recorded data. This also simplified bandpass optimisation, which was performed according to the instrumental bandpass rather than the bandpass after the action of the real-time dedispersion filter; and dedispersion, which was based directly on the retrospective GPS-derived STEC measurements, rather than the difference between these and the varying real-time dedispersion filter setting.

Analog-to-digital converters (ADCs) do not typically display perfectly linear behaviour. The non-linearity of the ADCs used in this experiment is known to cause errors of up to 2.5% in the recorded pulse amplitudes [21]. This is much less significant than the non-linearity observed earlier in the signal path (see Sec. 2.2), so it is not necessary to correct for it in real time. However, a small error in the pulse amplitude can lead to a large error in the number of recorded high-significance events, so it is desirable to remove the effect in retrospective processing. Since

the magnitude of the effect is different in each ADC, it causes differences in the number of high-significance events recorded in each beam, which we initially mistook for a real phenomenon rather than an instrumental effect. We calibrated the errors individually for each ADC and used these measurements to perform the linearity correction by rescaling the pulse amplitudes. For simplicity, we applied a correction averaged between positive and negative directions in voltage, neglecting the minor ($\pm 0.4\%$) variation between them.

Spectra from recorded events show common features from instrumental effects: at zero frequency, resulting from a direct-current offset in the ADCs; and at the maximum IF frequency (512 MHz), which is a harmonic of the operating frequency of the digital electronics. In real time, these were strongly suppressed by the dedispersion filter. For retrospective analysis, they were completely removed.

Although RFI is of greatest concern in this experiment when it is impulsive in nature and may be mistaken for an Askaryan pulse, continuous narrow-band RFI is also an unwanted source of noise. We dealt with this simply, removing any channels in the Fourier transforms of the data with a spectral density exceeding four times the RMS value. These constitute a smaller fraction of the band than in lower-frequency observations [e.g. 13] — typically 0–2 out of 2,041 channels — and their removal does not appreciably affect the sensitivity.

After the above steps and the bandpass optimisation from Sec. 5.1, the RMS signal power for each beam and polarisation was calculated and averaged over each 1-minute interval of observing time, to avoid the $\sim 2\%$ uncertainty associated with calculating this quantity from a single buffer of data while still measuring fluctuations in the instrumental gain over longer timescales. We then performed the remaining optimisation steps from Sec. 5, including dedispersion based on STEC values calculated once per minute of observing time. The magnitude of the most extreme peak in the data after this last step, divided by the RMS signal power, gives us the significance of the peak n_σ , as defined in Eq. 11, in units of σ .

6.2. Removal of RFI

Some RFI events escaped the real-time anticoincidence filter: although they featured peaks in multiple beams, the relative amplitudes of these peaks varied, so the triggering peak could exceed the trigger threshold in one beam while the corresponding peaks in the other beams remained below the threshold. These constituted a small fraction of the raw triggers, but were still numerous enough to dominate the population of recorded high-significance events, limiting the sensitivity of the experiment. To exclude them, we applied a series of cuts, tightening the exclusion criteria used in the real-time filter and discriminating against other features observed in the RFI events. As we have no expectation for the spectrum or dispersion of an RFI pulse, these cuts were applied to the raw data with no

dedispersion or bandpass optimisation; but they did include interpolation, formation of the signal envelope, and other corrections that were not performed in real time.

The initial *anticoincidence* cuts are a simple refinement of the real-time filter, exploiting the length of the buffers of recorded data (2–4 μs) and including the processing described in Sec. 6.1. They excluded

- #1: events with a peak in excess of 6.2σ in any beam other than the triggering beam; and
- #2: events with peaks in excess of 5.8σ in any two beams, or both polarisations of the same beam, other than the triggering beam,

which together constituted the majority of the recorded RFI. These thresholds were chosen as a compromise between discriminating power and the false exclusion rate: the probability that a genuine event will be misidentified as RFI and thus excluded. The multi-beam requirement in cut #2 allows the threshold to be set lower than in cut #1.

Visual inspection of RFI events, including those which passed cuts #1–2, frequently showed extended structure in the time domain: both broad pulses with a duration of several tens of ns and series of pulses with regular spacings of several hundreds of ns. We were able to experimentally reproduce the characteristic latter profile with the radio emission from an electric barbecue lighter, and we ascribe these events to other anthropogenic electrical activity with similar qualities. This motivated further cuts on the *width* of the pulses in a single beam in a single event, excluding

- #3: events with a peak in excess of 6.2σ in the same beam and polarisation as the peak which met the trigger condition, but separated from it by more than 10 ns; and
- #4: events with a peak in excess of 6.2σ in the same beam as the peak which met the trigger condition, but in the other polarisation, and separated from it by more than 60 ns.

These cuts discriminated both against repeated pulses and, to a lesser extent, broad single pulses. Cut #3 places an effective 10–20 ns upper limit on the width of a pulse for it to avoid being excluded as RFI, but this is consistent with the expected ~ 1 ns duration of an Askaryan pulse. The increased separation permitted in cut #4 is to allow for the remote possibility of a large error in the timing calibration between polarisations.

The cuts applied thus far have relatively lax thresholds and hence a low false exclusion rate, which we quantify with two techniques. The first is to determine the number of samples to which the exclusion threshold has been applied in each cut, and to determine the probability that it has been exceeded based on their expected Gaussian distribution and the consequent Rayleigh distribution of the signal envelope. The second is to operate our instrument with a minimal trigger threshold, continuously recording noise data without the usual bias towards high-amplitude

peaks, and subjecting this noise data to the same processing as the observational data. Applying the cuts to the noise data and measuring the fraction of excluded events gives a measure of the false exclusion rate which is sensitive to non-Gaussianity in the background noise, from both RFI and instrumental effects. Both approaches yield for cuts #1–4 a collective false exclusion rate of $< 0.1\%$, leading to a correspondingly insignificant loss of observing time.

The events excluded as RFI by cuts #1–4 are clustered in time on scales $\lesssim 10$ s (see Fig. 15), indicating that they originate from transient sources which are typically active for a similar period. The events which pass these cuts are significantly in excess of those expected from pure noise for peak amplitudes $\gtrsim 8\sigma$ (see Fig. 16), suggesting that they are dominated by RFI which escaped the cuts by not displaying the characteristics typical of RFI events: for example, having an origin in a sidelobe of one beam, but in nulls of the sidelobe patterns of the other beams, so that they do not detect coincident pulses. Based on this, we expect RFI events which passed the cuts to also display clustering behaviour with one another; but, as the positions of the sidelobes change over time (relative to local RFI sources, as the telescope tracks the Moon), they will generally occur at different times from RFI events which were already excluded, and not be clustered with them. This motivates a further cut based on the *proximity* of events to one another in time, in which we exclude

#5: events recorded within 10 s of another event with a peak exceeding 8σ , which also passed cuts #1–4.

We choose to apply an exclusion window only around events with peak amplitude $> 8\sigma$ because we expect these events to be dominated by RFI; and we ignore events excluded by previous cuts because, as argued above, we do not expect these to be clustered with RFI events remaining in our sample. In this way, we limit ourselves to exclusion windows of length 20 s centred on each of 727 events. These are, as expected, clustered, so the windows overlap and the total excluded time is only 1.1×10^4 s, equivalent to a 1.7% false exclusion rate.

The clustering exploited in cut #5 is too weak to act as a completely reliable discriminant against RFI, so we perform an additional cut to remove events with very weak coincident pulses in multiple beams. To minimise the false exclusion rate associated with a lower exclusion threshold, this *strict* cut searches only for very closely coincident pulses, excluding

#6: events with a 4.5σ peak within 20 ns of the position of the peak in the triggering beam, in another beam.

This ± 20 ns window is sufficiently wide to encompass the likely timing calibration error between beams, while minimising the probability that it will contain a peak exceeding the exclusion threshold purely from Gaussian noise. Nonetheless, due to the low threshold, this cut has the

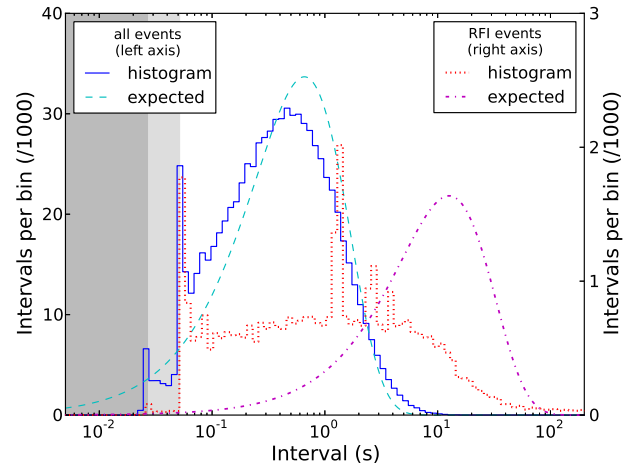


Figure 15: Distribution of intervals between events, for all recorded events (solid; left axis) and for events identified as RFI by cuts #1–4 (dotted; right axis) with a bin width of 0.05 decades. These two histograms have been slightly offset for display purposes. The expected distributions for a purely stochastic distribution of events are also shown. The distribution for all recorded events, which are dominated by thermal noise, fits quite closely to the expected distribution, with minor deviations caused by variation in the trigger rate during our observations. The events identified as RFI are typically separated by shorter intervals than expected, indicating that they are clustered in time. The shaded regions are excluded due to the maximum possible trigger rates for buffers of length $2 \mu\text{s}$ (dark) and $4 \mu\text{s}$ (light); the pile-up at their boundaries results from short periods when RFI or a decreased trigger threshold caused our instrument to trigger continuously. Peaks in the distribution for RFI events are visible at a period of 1.3 s and its subharmonics, indicating a source of RFI pulses with regular spacing of this width.

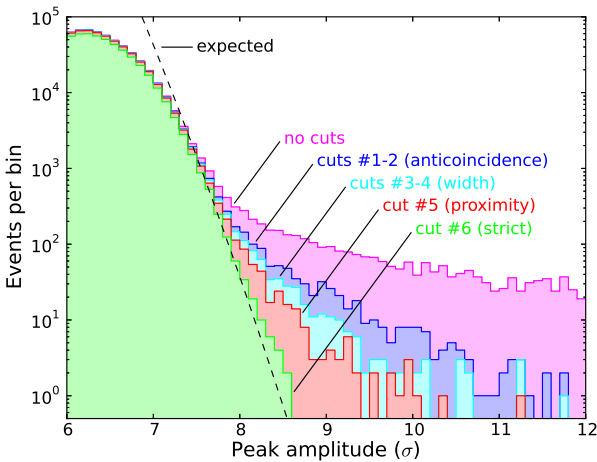


Figure 16: Distribution of amplitudes of events (bin width 0.1σ) for observing runs from May 2010 onwards, after each set of anti-RFI cuts except for cut #7, which has only a minimal effect (see Table 2). The expected distribution incorporates the effects of interpolation and formation of the signal envelope; the smearing effect due to random error from quantisation noise and variation in the RMS signal power [as per 21] and from the differential non-linearity between positive and negative voltage neglected in the earlier correction for ADC non-linearity; and the reduction in effective observing time described in Sec. 6.3.1, including the effects of the false exclusion rates for cuts #1–7. The deficit in the number of events at low amplitudes is due to these events not being consistently recorded in this experiment. The number of events with low amplitudes, which are dominated by thermal noise, is noticeably decreased by cut #6, due to its high false exclusion rate. After all cuts, the number of events at high amplitudes is consistent with expectations.

Table 2: Numbers of events remaining after each set of anti-RFI cuts.

Cut(s)	Criterion	Events	
		total	$> 8\sigma$
—	none	794 568	4 305
#1–2	anticoincidence	756 947	789
#3–4	width	755 867	483
#5	proximity	741 045	294
#6	strict	668 090	74
#7	spectrum	657 421	73

highest false exclusion rate: calculating this as for cuts #1–4, and conservatively taking the higher value produced by applying the cut to experimental noise data, the false exclusion rate is 5.4%. This cut also limits the sensitivity of this experiment due to the possibility of a lunar Askaryan pulse being seen in multiple beams and excluded as RFI.

Cuts #1–6 were sufficient to exclude impulsive RFI, but a further cut was necessary to identify periods of continuous narrow-band RFI which exceeded the dynamic range. A fluctuation in the voltage beyond the limits of the 8-bit digitisation range will be clipped to +127 ADU (if positive) or −128 ADU (if negative), and its full magnitude will be unrecorded. Powerful narrow-band RFI causes this clipping to happen frequently; although it is removed in retrospective processing, the problem occurs earlier, at the digitisation stage, so the information is already lost. We are therefore unable to reliably record high-significance events during periods of intense narrow-band RFI, so we remove these periods with a further cut on the *spectrum* of pulses, excluding

#7: events with more than 50% of the spectral power concentrated in less than 1% of the band, in any beam or polarisation.

Unlike previous cuts, this excludes very few high-significance events, but it does exclude all of the 24 events passing cuts #1–6 with sample values that reach the limits of the digitisation range, which may have been clipped. This ensures that any remaining high-significance events are recognised as such, rather than potentially being reconstructed with lower significance because their peak amplitude was reduced by clipping. The false exclusion rate, calculated as for previous cuts, is 0.7%.

Fig. 16 shows the distribution of the peak amplitudes after complete processing, including dedispersion and band-pass optimisation, and the effects on this distribution of the series of cuts described above. The numbers of events after each cut are also listed in Table 2. Events with amplitudes $> 8\sigma$ are initially dominated by RFI, but after all cuts there are only 73 such events against an expectation of 61 ± 8 . Of these, 3 events appear to have characteristics — broad, coincident or repeated pulses — suggesting that they may also be due to RFI, but the significance of these

features cannot be rigorously established through visual inspection, so we do not exclude these events. The remaining 70 events all appear to be consistent with the expected rare fluctuations in the background thermal noise.

The most significant event remaining after cuts has an amplitude of 8.6σ , with 3,000 more significant events being excluded. The real-time anticoincidence filter excluded raw triggers at a ratio of $\sim 150 : 1$ (see Table 1). As the amplitudes of events excluded in real time are not recorded, we cannot state with certainty the total number of excluded events with amplitudes exceeding the 8.6σ significance threshold; but, if the distribution of amplitudes is the same for raw triggers as for recorded events, then the combination of the real-time filter and subsequent cuts has excluded $\sim 450,000$ high-significance RFI events and accepted none. Since raw triggers are dominated by RFI, we expect them to be biased towards larger amplitudes, in which case the number of successfully rejected high-significance events is even greater than this.

6.3. Threshold

The 8.6σ peak amplitude of the most significant event remaining after cuts defines the significance threshold of this experiment. If an event had been detected in excess of this threshold, the confidence of its identification as a lunar-origin Askaryan pulse would not be rigorously established, as the cuts were not all defined *a priori*. However, the absence of such events allows us to place a limit on the intensity of radio pulses originating from the Moon during this experiment; which, combined with the rigorous testing of the false exclusion rates of the cuts to determine the effective observing time, allows limits to be placed on the fluxes of UHE cosmic rays and neutrinos.

Thus far, however, we have only established that there are no events exceeding this threshold in the recorded data. We must also establish that any such event occurring during this experiment would have been recorded in the first place. Neglecting effects which merely reduce the effective observing time, which we discuss in Sec. 6.3.1, there are two possible cases in which we could fail to record a high-significance event.

In the first case, a pulse might be properly sampled and processed, but its peak amplitude imperfectly reconstructed by the suboptimal real-time processing, so that it falls below the trigger threshold and is not recorded. The maximum loss of amplitude in real-time reconstruction relative to full retrospective processing is 23%, as found in Sec. 5.5. For a peak with a full magnitude of 8.6σ or greater, this means that it will be detected in real time with a minimum magnitude of 6.6σ . Since this is still above the 6.4σ trigger threshold, all such events are recorded and will be correctly analysed retrospectively, so this case does not occur.

In the second case, a pulse might exceed the $\pm 12\sigma$ digitisation range of the ADCs, causing it to be clipped and its full magnitude to go unrecorded, with its amplitude then further decreased by the real-time dedispersion filter to a

level below the trigger threshold. While it is not generally expected for dedispersion to reduce the amplitude of a dispersed lunar-origin Askaryan pulse as required in this case, note that both clipping and the surrounding thermal noise may alter the inherent dispersion of the pulse, and the real-time dedispersion filter may change both the phase of the pulse and the phase of the sampling times relative to its peak, with consequent effects on its amplitude. Taking this reduction in amplitude, for the maximum setting of the filter, to be equal to the expected loss from dispersion for the maximum ionospheric STEC during our observations, and including the effects of partial interpolation, the maximum signal loss is 32%, found with the same simulation used in Sec. 5.5. This is sufficient to reduce the amplitude of a clipped pulse from 12σ to 8.2σ . As this is still above the 6.4σ trigger threshold, all such events are recorded, and this case also does not occur; so the absence of any clipped pulses after cut #7 implies that clipped events occurred only as RFI.

With these possibilities eliminated, we have established that no lunar-origin pulses occurred during our experiment (within the effective observing time; see Sec. 6.3.1) with a reconstructed pulse height in excess of 8.6σ . This threshold does not allow for loss of amplitude due to suboptimality of the final reconstruction ($< 0.4\%$; see Sec. 5.5), or from Faraday rotation ($\lesssim 1\%$; see Sec. 2.1), but these are small compared to the calibration uncertainty of $\pm 4.5\%$ (random) $\pm 4\%$ (systematic) found for the electric field in Sec. 3. We convert this threshold significance $n_\sigma = 8.6$ (see Eq. 11) into a threshold spectral electric field strength with Eq. 10, using the calibrated sensitivity for each beam pointing and polarisation.

For a limb beam pointing (see Fig. 3), this gives a threshold of $0.0047 \mu\text{V}/\text{m}/\text{MHz}$ for a pulse originating at the centre of the beam, polarised radially to the Moon. At the closest point on the lunar limb, the threshold is $0.0053 \mu\text{V}/\text{m}/\text{MHz}$ for the same polarisation. Further away on the limb, the sensitivity to a radially-polarised pulse decreases, both due to the greater distance from the centre of the beam and due to misalignment between the polarisation of the beam and that of the pulse. The fraction of the lunar limb over which the threshold takes no more than $\sqrt{2}$ times its minimum value on the limb (i.e. at least half sensitivity in the power domain) is 10% for a single limb beam.

For a half-limb beam pointing, the threshold is $0.0074 \mu\text{V}/\text{m}/\text{MHz}$ for either polarisation (both at 45° to a radial alignment) at the centre of the beam, or $0.010 \mu\text{V}/\text{m}/\text{MHz}$ at the closest point on the limb. This is equivalent to a threshold of $0.014 \mu\text{V}/\text{m}/\text{MHz}$ for a radial pulse at this point. Further away on the limb, a radially aligned pulse more closely matches one or the other polarisation of the beam, which results in greater limb coverage: 20%, calculated as for the limb beam. These limb coverage values are suitable for inclusion in particle aperture models such as that of Gayley et al. [39]. The limb coverage for both beams is shown in Fig. 17.

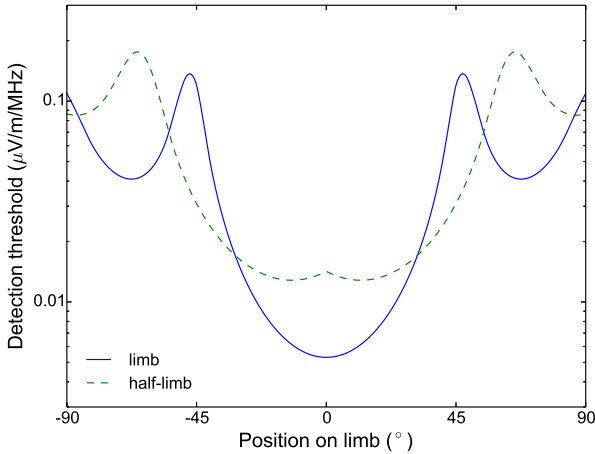


Figure 17: Detection threshold for a radially-aligned pulse along the lunar limb for the limb and half-limb beams (see Fig. 3). A limb position of 0° corresponds to a pulse originating at the point on the limb closest to the beam. The limb beam is most sensitive, but the sensitivity declines rapidly for a pulse away from this point, which is both more distant from the centre of the beam and suboptimally aligned for the radial polarisation channel. The half-limb beam is less sensitive, but a pulse slightly removed from the closest point on the limb will be better aligned for one or the other polarisation channel, resulting in greater limb coverage.

For comparison, the two previous lunar radio experiments with the lowest detection thresholds are LUNASKA ATCA with a threshold of $0.0145\text{--}0.016\text{ }\mu\text{V/m/MHz}$ [11] and RESUN with a threshold of $0.017\text{ }\mu\text{V/m/MHz}$ [14]. (Jaeger et al. [14] report a threshold of $0.013\text{ }\mu\text{V/m/MHz}$ for the Kalyazin experiment [10], but this is not compatible with the originally published flux density threshold of 13.5 kJy .) Consequently, the threshold for the radial polarisation of our limb beams improves over previous experiments by a factor of three. This results in a decrease by a corresponding factor in the minimum detectable particle energy compared to these previous experiments.

Note that each of these thresholds represents the spectral electric field strength for which a pulse has a 50% probability of detection: the addition of thermal noise may cause pulses over the threshold to be missed, or pulses under the threshold to be detected. The amplitude of the noise is $1/8.6 \sim 12\%$ of the threshold value. The effect of thermal noise on the aperture for UHE particles is incorporated into simulations such as those of James and Protheroe [5], although Gayley et al. [39] argue that the effect is negligible.

6.3.1. Effective observing time

The effective observing time is decreased from the total time reported in Table 1, either through a decrease in the duty cycle or a chance of incorrectly excluding a detected Askaryan pulse, by three effects.

- The real-time anticoincidence filter excludes events in which the trigger threshold is exceeded in multiple

Table 3: Effective duration of observations, after losses due to the real-time anticoincidence filter, time taken for data storage, and false exclusion rates of anti-RFI cuts.

Configurations	A, B ^a	A*, B* ^b
Total duration	116.9 hrs	31.8 hrs
Real-time anticoincidence	−0.4%	−0.2%
Data storage	−6.8%	−6.1%
Anti-RFI cuts	−8.4%	−6.7%
Effective duration	99.4 hrs	27.8 hrs

^a With two limb beams and one half-limb beam.

^b With two limb beams.

beams, as described in Sec. 2. It applies a 200 ns exclusion window around each trigger, with the trigger rate being dominated by the off-Moon beam. The number of such triggers is logged, so an upper limit (assuming no overlapping windows) to the total excluded time is known.

- After a successful trigger, the system is temporarily unable to respond to a further trigger as it copies the previous event to permanent storage, as described in Sec. 2. We measured the length of this interval by operating briefly with a minimal trigger threshold to cause continuous triggering, and noting the number of recorded events. The inverse of this event rate gives us the excluded interval after each recorded event, which is 27 ms when the buffer length is set to $2\text{ }\mu\text{s}$ and 51 ms for a buffer length of $4\text{ }\mu\text{s}$.
- Each anti-RFI cut has a corresponding false exclusion rate as described in Sec. 6.2, evaluated primarily by applying the same cuts to a sample of background noise data.

We consider the above effects separately for observing time with two on-Moon beams (the September observing run) and three on-Moon beams (all other observing runs), with the results shown in Table 3. RFI was more prevalent in the September observing run, but was more effectively filtered out by our initial anticoincidence and width cuts, probably due to the use of an additional off-Moon beam, so less time was lost to the proximity cut than in other observing runs. The total effective observing time, combining all observing runs, is 127.2 hours with each of two limb beams and 99.4 hours with one half-limb beam.

7. Conclusion and summary

We have conducted an experiment to search for radio pulses from particle cascades initiated by UHE particles interacting in the Moon. Our anticoincidence filter and subsequent cuts exclude effectively all RFI, and the remaining events are consistent with those expected from thermal noise. We detected no pulses originating on the limb of the Moon with an electric field exceeding $0.0053\text{ }\mu\text{V/m/MHz}$,

assuming linear polarisation radial to the Moon, during 127 hours of effective observing time. This threshold is an improvement over previous lunar radio experiments by a factor of three. Our non-detection implies limits on the fluxes of UHE cosmic rays and neutrinos, which may be modeled based on the parameters reported here.

We have achieved, for the first time, effectively complete exclusion of RFI pulses without the benefit of a coincidence requirement between multiple antennas or bands. Previous experiments have either required a coincidence between multiple channels, preventing them from achieving the full coherent sensitivity possible with their combined bandwidth and collecting area, or been dominated by remnant RFI.

7.1. Considerations for future experiments

Future lunar radio experiments will take advantage of new and improved radio telescopes which are currently under development, achieving greater sensitivity to UHE particles. In particular, Singh et al. [40] have proposed lunar radio observations with LOFAR, an aperture array telescope currently in its commissioning phase, and Bray et al. [41] have proposed continued use of the Parkes radio telescope used in this experiment, using one of the Phased Array Feed (PAF) receivers recently developed for ASKAP [42]. Both of these proposed experiments would utilise multiple beams pointing on the Moon with a real-time anticoincidence filtering scheme to identify lunar-origin Askaryan pulses. Here we summarise the experimental considerations addressed in this work that may be relevant for these and other future experiments.

The optimum signal path for a lunar radio experiment can be understood in terms of the aspects of a matched filter, with the phase, dispersion and bandpass optimised to match the expected characteristics of an Askaryan pulse. In each of these respects, this experiment improves on previous work.

Phase: In typical radio receivers, the signal is shifted to an intermediate frequency by mixing it with a local oscillator signal, effectively randomising the phase of a pulse and potentially decreasing its peak amplitude (by up to 18%, in this case). We form the signal envelope to restore the maximum peak amplitude, but this also increases the noise level, partly offsetting the improvement in the signal-to-noise ratio. In some recent radio telescopes the radio frequency signal is sampled directly; provided that it is Nyquist-sampled (i.e. no aliasing occurs), this preserves the original phase of a pulse. The inherent phase of an Askaryan pulse corresponds to the worst-case peak amplitude; but, in this case, the maximum peak amplitude can be restored by performing a Hilbert transform on the signal, with no change in the noise level.

Dispersion: Our experiment is the first to use a dedispersion filter which was adjusted in real time based

on changes in the ionospheric STEC along the line of sight to the Moon. The precision of this correction is limited by the availability of precise STEC measurements, particularly in real time. The most precise values available are from GPS data, with the uncertainty dominated by the conversion from STEC (between ground station and GPS satellite) to VTEC and back. Future experiments, which will generally be more affected by dispersion due to increased bandwidth, could directly use GPS STEC values measured from a ground station located at the telescope site, or measure the STEC through Faraday rotation of polarised lunar thermal emission as proposed by McFadden et al. [43].

Bandpass: For this experiment, we optimised the bandpass for sensitivity to lower-energy particles. The improvement in sensitivity in this case is an insignificant 1–2%, but will become increasingly significant for future experiments with larger fractional bandwidths.

The loss of sensitivity due to a finite sampling rate can be negated by interpolating to reconstruct intermediate sample values, as first suggested in this context by James et al. [11]. For a triggered experiment, this interpolation may need to be at least partially implemented in real time, as in this work. The potential loss of sensitivity without this correction depends on the oversampling factor, and is 22% for our experiment.

An improvement in sensitivity is obtained by aligning a linearly-polarised receiver with the expected polarisation of an Askaryan pulse, at least when observing at high radio frequencies for which Faraday rotation is minimal. For future experiments, with sufficiently precise calibration of the delay and phase between polarisation channels of each beam, it may be possible to construct signals corresponding to arbitrary linear polarisations, allowing this benefit to be obtained regardless of the orientation and native polarisation of the receiver.

An understanding of the above considerations is required in order to properly evaluate the sensitivity of a lunar radio experiment to Askaryan pulses. We have described techniques which allow this sensitivity to be maximised. The implementation of these techniques in future lunar radio experiments, and eventually with the SKA, will result in the greatest possible aperture for the detection of UHE particles.

Acknowledgements

The Parkes radio telescope is part of the Australia Telescope which is funded by the Commonwealth of Australia for operation as a National Facility managed by CSIRO. Some data used in this study were acquired as part of NASA’s Earth Science Data Systems and archived and distributed by the Crustal Dynamics Data Information System (CDDIS). This research was supported by the

Australian Research Council's Discovery Project funding scheme (project number DP0881006). JDB acknowledges support from ERC-StG 307215 (LODESTONE).

References

- [1] Pierre Auger Collaboration, Properties and performance of the prototype instrument for the Pierre Auger Observatory, *Nucl. Instrum. Meth. A* 523 (2004) 50–95.
- [2] Telescope Array Collaboration, The surface detector array of the Telescope Array experiment, *Nucl. Instrum. Meth. A* 689 (2012) 87–97.
- [3] R. D. Dagkesamanskii, I. M. Zheleznykh, Radio-astronomy method for detecting neutrinos and other elementary particles of superhigh energy, *Sov. Phys. JETP Lett.* 50 (1989) 259–261.
- [4] G. A. Askaryan, Excess Negative Charge of an Electron-Photon Shower and its Coherent Radio Emission, *Sov. Phys. JETP* 14 (2) (1962) 441–443.
- [5] C. W. James, R. J. Protheroe, The sensitivity of the next generation of lunar Cherenkov observations to UHE neutrinos and cosmic rays, *Astropart. Phys.* 30 (2009) 318–332.
- [6] S. ter Veen, S. Buitink, H. Falcke, C. W. James, M. Mevius, O. Scholten, K. Singh, B. Stappers, K. D. de Vries, Limit on the ultrahigh-energy cosmic-ray flux with the Westerbork synthesis radio telescope, *Phys. Rev. D* 82 (10) (2010) 103014.
- [7] Y. S. Jeong, M. H. Reno, I. Sarcevic, Radio Cherenkov signals from the Moon: Neutrinos and cosmic rays, *Astropart. Phys.* 35 (2012) 383–395.
- [8] T. H. Hankins, R. D. Ekers, J. D. O'Sullivan, A search for lunar radio Čerenkov emission from high-energy neutrinos, *MNRAS* 283 (1996) 1027–1030.
- [9] P. W. Gorham, C. L. Hebert, K. M. Liewer, C. J. Naudet, D. Saltzberg, D. Williams, Experimental Limit on the Cosmic Diffuse Ultrahigh Energy Neutrino Flux, *Phys. Rev. Lett.* 93 (4) (2004) 041101.
- [10] A. R. Beresnyak, R. D. Dagkesamanskii, I. M. Zheleznykh, A. V. Kovalenko, V. V. Oreshko, Limits on the Flux of Ultrahigh-Energy Neutrinos from Radio Astronomical Observations, *Astron. Rep.* 49 (2005) 127–133.
- [11] C. W. James, R. D. Ekers, J. Alvarez-Muñiz, J. D. Bray, R. A. McFadden, C. J. Phillips, R. J. Protheroe, P. Roberts, LUNASKA experiments using the Australia Telescope Compact Array to search for ultrahigh energy neutrinos and develop technology for the lunar Cherenkov technique, *Phys. Rev. D* 81 (4) (2010) 042003.
- [12] R. E. Spencer, A. Macfarlane, O. Mills, L. Piccirillo, La Luna: Lovell Attempts LUNAR Neutrino Acquisition, in: *Proc. EVN Symp.* 2010, *Proc. Sci.*, 097, 2010.
- [13] S. Buitink, O. Scholten, J. Bacelar, R. Braun, A. G. de Bruyn, H. Falcke, K. Singh, B. Stappers, R. G. Strom, R. A. Yahyaoui, Constraints on the flux of ultra-high energy neutrinos from Westerbork Synthesis Radio Telescope observations, *A&A* 521 (2010) A47.
- [14] T. R. Jaeger, R. L. Mutel, K. G. Gayley, Project RESUN, a Radio EVLA Search for UHE Neutrinos, *Astropart. Phys.* 34 (2010) 293–303.
- [15] C. L. Carilli, S. Rawlings, Motivation, key science projects, standards and assumptions, *New Astron. Rev.* 48 (2004) 979–984.
- [16] K. Greisen, End to the Cosmic-Ray Spectrum?, *Phys. Rev. Lett.* 16 (1966) 748–750.
- [17] G. T. Zatsepin, V. A. Kuzmin, Upper Limit of the Spectrum of Cosmic Rays, *Sov. Phys. JETP Lett.* 4 (1966) 78.
- [18] R. D. Ekers, C. W. James, R. J. Protheroe, R. A. McFadden, Lunar radio Cherenkov observations of UHE neutrinos, *Nucl. Instrum. Meth. A* 604 (2009) 106.
- [19] J. D. Bray, R. D. Ekers, P. Roberts, J. E. Reynolds, C. W. James, C. J. Phillips, R. J. Protheroe, R. A. McFadden, M. G. Aartsen, A limit on the ultra-high-energy neutrino flux from lunar observations with the Parkes radio telescope, submitted, 2014.
- [20] L. Staveley-Smith, W. E. Wilson, T. S. Bird, M. J. Disney, R. D. Ekers, K. C. Freeman, R. F. Haynes, M. W. Sinclair, R. A. Vaile, R. L. Webster, A. E. Wright, The Parkes 21 cm multibeam receiver, *PASA* 13 (1996) 243–248.
- [21] J. D. Bray, R. D. Ekers, P. Roberts, Noise statistics in a fast digital radio receiver: the Bedlam backend for the Parkes radio telescope, *Exp. Astron.* 36 (1–2) (2013) 155–174.
- [22] C. W. James, R. J. Protheroe, The directional dependence of apertures, limits and sensitivity of the lunar Cherenkov technique to a UHE neutrino flux, *Astropart. Phys.* 31 (2009) 392–398.
- [23] Pierre Auger Collaboration, Update on the correlation of the highest energy cosmic rays with nearby extragalactic matter, *Astropart. Phys.* 34 (2010) 314–326.
- [24] A. Savitzky, M. J. E. Golay, Smoothing and Differentiation of Data by Simplified Least Squares Procedures., *Analytical Chem.* 36 (8) (1964) 1627–1639.
- [25] P. H. Moffat, Aperture synthesis polarimetry of the Moon at 21 cm, *MNRAS* 160 (1972) 139–154.
- [26] V. S. Troitskii, T. V. Tikhonova, Thermal radiation from the moon and the physical properties of the upper lunar layer, *Rad. & Quan. Elec.* 13 (1970) 981–1010.
- [27] P. Roberts, Ionospheric dispersion compensation using a novel microwave de-dispersion filter, in: *Proc. WARS 2008*, arXiv:1305.0333, 2008.
- [28] M. Hernández-Pajares, J. Juan, J. Sanz, R. Orus, A. Garcia-Rigo, J. Feltens, A. Komjathy, S. Schaer, A. Krankowski, The IGS VTEC maps: a reliable source of ionospheric information since 1998, *J. Geodesy* 83 (2009) 263–275.
- [29] C. Noll, The crustal dynamics data information system: A resource to support scientific analysis using space geodesy, *Adv. Space Res.* 45 (12) (2010) 1421–1440.
- [30] S. Schaer, W. Gurtner, J. Feltens, IONEX: The IONosphere Map EXchange Format Version 1, in: *Proc. IGS Analysis Center Workshop 1998*, 233–247, 1998.
- [31] D. L. Gallagher, P. D. Craven, R. H. Comfort, An empirical model of the earth's plasmasphere, *Adv. Space Res.* 8 (1988) 15–21.
- [32] C. L. Byrne, Signal processing: a mathematical approach, A. K. Peters, ISBN 1-56881-242-6, 2005.
- [33] J. Alvarez-Muñiz, E. Marqués, R. A. Vázquez, E. Zas, Coherent radio pulses from showers in different media: A unified parametrization, *Phys. Rev. D* 74 (2) (2006) 023007.
- [34] P. Miočinić, R. C. Field, P. W. Gorham, E. Guillian, R. Milinčić, D. Saltzberg, D. Walz, D. Williams, Time-domain measurement of broadband coherent Cherenkov radiation, *Phys. Rev. D* 74 (4) (2006) 043002.
- [35] J. Alvarez-Muñiz, A. Romero-Wolf, E. Zas, Čerenkov radio pulses from electromagnetic showers in the time domain, *Phys. Rev. D* 81 (12) (2010) 123009.
- [36] J. Alvarez-Muñiz, R. A. Vázquez, E. Zas, Calculation methods for radio pulses from high energy showers, *Phys. Rev. D* 62 (6) (2000) 063001.
- [37] M. S. Longuet-Higgins, The Statistical Analysis of a Random Moving Surface, *Phil. Trans. R. Soc. A* 249 (1957) 321–387.
- [38] T. K. Gaisser, A. M. Hillas, Reliability of the method of constant intensity cuts for reconstructing the average development of vertical showers, in: *Proc. ICRC 1977*, vol. 8, 353–357, 1977.
- [39] K. G. Gayley, R. L. Mutel, T. R. Jaeger, Analytic Aperture Calculation and Scaling Laws for Radio Detection of Lunar-target Ultra-high-energy Neutrinos, *ApJ* 706 (2009) 1556–1570.
- [40] K. Singh, M. Mevius, O. Scholten, J. M. Anderson, A. van Ardenne, M. Arts, M. Avruha, A. Asgekar, M. Bell, P. Benna, et al., Optimized trigger for ultra-high-energy cosmic-ray and neutrino observations with the low frequency radio array, *Nucl. Instrum. Meth. A* 664 (2012) 171–185.
- [41] J. D. Bray, R. D. Ekers, R. J. Protheroe, C. W. James, C. J. Phillips, P. Roberts, A. Brown, J. E. Reynolds, R. A. McFadden, M. Aartsen, LUNASKA neutrino search with the Parkes and ATCA telescopes, in: *Proc. ARENA 2012*, vol. 1535 of *AIP Conf. Series*, 21–26, 2013.

- [42] A. E. Schinckel, J. D. Bunton, T. J. Cornwell, I. Feain, S. G. Hay, The Australian SKA Pathfinder, in: Ground-based and Airborne Telescopes IV, vol. 8444 of *Proc. SPIE*, 84442A, 2012.
- [43] R. A. McFadden, R. D. Ekers, J. D. Bray, Ionospheric propagation effects for UHE neutrino detection with the lunar Cherenkov technique, in: Proc. ICRC 2011, vol. 4, 284, 2011.